



Junos[®] OS

Broadband Subscriber Access Protocols Feature Guide



Modified: 2018-03-08

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be 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® OS Broadband Subscriber Access Protocols Feature Guide
Copyright © 2018 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 <https://www.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xxiii
	Documentation and Release Notes	xxiii
	Supported Platforms	xxiii
	Using the Examples in This Manual	xxiii
	Merging a Full Example	xxiv
	Merging a Snippet	xxiv
	Documentation Conventions	xxv
	Documentation Feedback	xxvii
	Requesting Technical Support	xxvii
	Self-Help Online Tools and Resources	xxvii
	Opening a Case with JTAC	xxviii
Part 1	Broadband Subscriber Access Network Overview	
Chapter 1	Broadband Subscriber Access Network Overview	3
	Subscriber Access Network Overview	3
	Multiservice Access Node Overview	4
	LDP Pseudowire Autosensing Overview	5
	Pseudowire Ingress Termination Background	6
	Pseudowire Autosensing Approach	7
	Sample Configuration	8
	Layer 2 Services on Pseudowire Service Interface Overview	8
	Traffic from Customer LAN to MPLS	9
	Traffic from Service Edge to Customer LAN	10
	Pseudowire Service Interfaces	11
	Sample Configuration	11
	Ethernet MSAN Aggregation Options	15
	Direct Connection	16
	Ethernet Aggregation Switch Connection	16
	Ring Aggregation Connection	17
	Broadband Access Service Delivery Options	17
	Digital Subscriber Line	17
	Active Ethernet	18
	Passive Optical Networking	18
	Hybrid Fiber Coaxial	18
	Broadband Delivery and FTTx	19

Part 2	Configuring the DHCP Access Network	
Chapter 2	Configuring Services for DHCP Subscribers	23
	DHCP and Subscriber Management Overview	23
	Extended DHCP Local Server and Subscriber Management Overview	23
	Extended DHCP Relay and Subscriber Management Overview	24
	DHCP Relay Proxy and Subscriber Management Overview	24
	Subscriber Access Operation Flow Using DHCP Relay	24
	Defining Various Levels of Services for DHCP Subscribers	25
	Example: Configuring a Tiered Service Profile for Subscriber Access	26
Chapter 3	Applying RADIUS Route Attributes to Subscribers or to Access Networks	31
	Access and Access-Internal Routes for Subscriber Management	31
	Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management	32
	Configuring Dynamic Access Routes for Subscriber Management	33
Chapter 4	Suppressing DHCP Access, Access-Internal, and Destination Routes	37
	Suppressing DHCP Access, Access-Internal, and Destination Routes	37
	Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default	38
Chapter 5	Providing Security in the DHCP Network	41
	DHCP Snooping Support	41
	Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server	43
	Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent	44
	Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent	48
	Disabling DHCP Snooping Filters	51
	Example: Configuring DHCP Snooping Support for DHCP Relay Agent	52
	Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent	54
	Preventing DHCP Spoofing	58
Chapter 6	Distinguishing Between Duplicate DHCPv4 Subscribers on the Same Subnet	61
	DHCPv4 Duplicate Client In Subnet Overview	61
	Guidelines for Configuring Support for DHCPv4 Duplicate Clients	62
	Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information	63
	Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces	64
Chapter 7	Distinguishing Between Duplicate DHCPv6 Subscribers	67
	DHCPv6 Duplicate Client DUIDs	67
	Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs	68

Chapter 8	Using the DHCP Relay Agent to Selectively Process DHCP Client Traffic	71
	DHCP Options and Selective Traffic Processing Overview	71
	Using DHCP Option Information to Selectively Process DHCP Client Traffic	73
	Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings	75
	Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing	79
Chapter 9	Configuring High Availability in the DHCP Access Network	85
	DHCP Liveness Detection Overview	85
	Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD	87
	Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients	89
	Configuring Detection of DHCP Local Server Client Connectivity with BFD	92
	Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients	94
	DHCP Liveness Detection Using ARP and Neighbor Discovery Packets	98
	How DHCP Liveness Detection with ARP and Neighbor Discovery Packets Works	98
	Send Functionality	98
	Receive Functionality	100
	Configuring BNG Detection of DHCP Local Server Client Connectivity with ARP and ND Packets	101
	Configuring BNG Detection of DHCP Relay Client Connectivity with ARP and ND Packets	103
	Configuring DHCP Host Detection of Client Connectivity with ARP and ND Packets	105
	High Availability Using Unified ISSU in the DHCP Access Network	106
	Graceful Routing Engine Switchover for DHCP	106
	Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover	107
	Benefits of Delaying Removal of Access Routes and Access-Internal Routes	107
	Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes	108
	Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes	108
	Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover	108
Chapter 10	Monitoring and Managing DHCP for Subscriber Access	111
	Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers	111
	Verifying and Monitoring Subscriber Management Unified ISSU State	111

Part 3	Configuring the PPP Access Network	
Chapter 11	Configuring PPP for Subscriber Access	115
	Dynamic Profiles for PPP Subscriber Interfaces Overview	115
	Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests	116
	Benefits of PPP Fast Keepalives	116
	How PPP Fast Keepalive Processing Works	116
	Statistics Display for PPP Fast Keepalive	117
	Effect of Changing the Forwarding Class Configuration	117
	Preventing the Validation of PPP Magic Numbers	117
	Configuring Dynamic Profiles for PPP	118
	Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges	120
	Attaching Dynamic Profiles to Static PPP Subscriber Interfaces	121
	Example: Minimum PPPoE Dynamic Profile	121
Chapter 12	Applying RADIUS Route Attributes to Subscribers or Access Networks	123
	Configuring Dynamic Access-Internal Routes for PPP Subscriber Management	123
	Verifying the Configuration of Access and Access-Internal Routes for PPP Subscribers	124
Chapter 13	Configuring Authentication for PPP	125
	Configuring Dynamic Authentication for PPP Subscribers	125
	Modifying the CHAP Challenge Length	127
Chapter 14	Configuring PPP Network Control Protocol Negotiation	129
	PPP Network Control Protocol Negotiation Mode Overview	129
	PPP NCP Negotiation Modes	129
	PPP NCP Negotiation Mode Supported Configurations	130
	PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers	130
	PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers	131
	PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations	131
	Controlling the Negotiation Order of PPP Authentication Protocols	132
	Configuring the PPP Network Control Protocol Negotiation Mode	134
	Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses	136
Chapter 15	Configuring High Availability in the PPP Access Network	139
	High Availability Using Unified ISSU in the PPP Access Network	139
	Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover	140
	Benefits of Delaying Removal of Access Routes and Access-Internal Routes	140
	Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes	140

	Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes	140
	Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover	141
Chapter 16	Monitoring and Managing PPP for Subscriber Access	143
	Verifying and Managing PPP Configuration for Subscriber Management	143
	Verifying and Monitoring Subscriber Management Unified ISSU State	143
Part 4	Configuring the L2TP Access Network	
Chapter 17	L2TP and Subscriber Access Overview	147
	L2TP for Subscriber Access Overview	147
	L2TP Terminology	149
	L2TP Implementation	150
	Sequence of Events on the LAC	151
	Sequence of Events on the LNS	152
Chapter 18	Configuring L2TP Tunneling and Switching for Subscribers	153
	L2TP Tunnel Switching Overview	153
	Application of Tunnel Switch Profiles	155
	Termination of Tunnel-Switched Sessions on the LTS	155
	Tunnel Switching Actions for L2TP AVPs at the Switching Boundary	157
	Configuring L2TP Tunnel Switching	161
	Setting the L2TP Receive Window Size	163
	Setting the L2TP Tunnel Idle Timeout	163
	Setting the L2TP Destruct Timeout	164
	Configuring the L2TP Destination Lockout Timeout	165
	Removing an L2TP Destination from the Destination Lockout List	165
	Configuring L2TP Drain	166
	Using the Same L2TP Tunnel for Injection and Duplication of IP Packets	167
Chapter 19	Configuring L2TP Control Messages for Subscribers	169
	Retransmission of L2TP Control Messages	169
	Configuring Retransmission Attributes for L2TP Control Messages	171
Chapter 20	Configuring L2TP LAC Subscribers	173
	Configuring an L2TP LAC	173
	Configuring How the LAC Responds to Address and Port Changes Requested by the LNS	175
	LAC Interoperation with Third-Party LNS Devices	177
	Globally Configuring the LAC to Interoperate with Cisco LNS Devices	178
	Configuring Username Modification for Subscriber Sessions	179
Chapter 21	Configuring L2TP LAC Tunneling for Subscribers	183
	LAC Tunnel Selection Overview	183
	Selection When Failover Between Preference Levels Is Configured	186
	Selection When Failover Within a Preference Level Is Configured	191

	Selection When Distributing the Session Load Across Multiple LNSs	193
	Weighted Load Balancing	193
	Destination-Equal Load Balancing	194
	L2TP Session Limits Overview	199
	Scenario 1: Chassis Limit	200
	Scenario 2: Tunnel Limit	200
	Scenario 3: Tunnel Group Limit	201
	Scenario 4: Session-Limit Group Limit	201
	Scenario 5: Individual Client Limit	203
	Limiting the Number of L2TP Sessions Allowed by the LAC or LNS	204
	Setting the Format for the Tunnel Name	206
	Configuring a Tunnel Profile for Subscriber Access	207
	Configuring the L2TP LAC Tunnel Selection Parameters	209
	Configuring LAC Tunnel Selection Failover Within a Preference Level	210
	Configuring Weighted Load Balancing for LAC Tunnel Sessions	211
	Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions	211
	Enabling the LAC for IPv6 Services	212
Chapter 22	Configuring Use of Subscriber Access Line and Connect Speed Information	213
	Subscriber Access Line Information Handling by the LAC and LNS Overview . . .	213
	Access Line Information Forwarding	213
	Access Line Information AVPs	214
	Connection Speed Updates on the LAC	216
	Connection Speed Updates on the LNS	217
	Interaction Between Global and Per-Destination Configurations	217
	Transmission of Tx and Rx Connection Speeds from LAC to LNS	219
	Methods for Determining the Speed Values Reported to the LNS	219
	Determining Initial Connect Speeds	223
	Fallback Mechanism for Connect Speed Values	224
	Transmission of the Receive Connect Speed AVP When Transmit and Receive	
	Connect Speeds are Equal	227
	Configuring the Method to Derive the LAC Connection Speeds Sent to the	
	LNS	228
	Configuring the Reporting and Processing of Subscriber Access Line	
	Information	230
	Preventing the LAC from Sending Calling Number AVP 22 to the LNS	235
	Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds	235
Chapter 23	Configuring L2TP LNS Inline Service Interfaces	239
	Configuring an L2TP LNS with Inline Service Interfaces	239
	Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface . .	242
	Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile . .	245
	Configuring an L2TP Access Profile on the LNS	246
	Configuring a AAA Local Access Profile on the LNS	248
	Configuring an Address-Assignment Pool for L2TP LNS with Inline Services . .	249
	Configuring the L2TP LNS Peer Interface	251
	Enabling Inline Service Interfaces	251
	Configuring an Inline Service Interface for L2TP LNS	252

	Configuring Options for the LNS Inline Services Logical Interface	253
	LNS 1:1 Stateful Redundancy Overview	254
	Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces	255
	Verifying LNS Aggregated Inline Service Interface 1:1 Redundancy	257
	L2TP Session Limits and Load Balancing for Service Interfaces	259
	Session Limits on Service Interfaces	259
	Session Load Balancing Across Service Interfaces	260
	Example: Configuring an L2TP LNS	263
	Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces	274
	Applying Services to an L2TP Session Without Using RADIUS	277
	Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions	284
	Configuring a Dynamic Profile for Dynamic LNS Sessions	285
Chapter 24	Configuring IP Packet Fragment Reassembly	289
	IP Packet Fragment Reassembly for L2TP Overview	289
	Configuring IP Inline Reassembly for L2TP	290
Chapter 25	Configuring High Availability in the L2TP Access Network	293
	L2TP and Graceful Routing Engine Switchover	293
	L2TP Failover and Peer Resynchronization	294
	Configuring the L2TP Peer Resynchronization Method	295
	High Availability Using Unified ISSU in the L2TP Access Network	297
	Verifying and Monitoring Subscriber Management Unified ISSU State	298
Chapter 26	Monitoring and Managing L2TP for Subscriber Access	299
	Verifying and Managing L2TP for Subscriber Access	299
	Testing L2TP Tunnel Configurations from the LAC	300
	Enabling Tunnel and Global Counters for SNMP Statistics Collection	303
Part 5	Configuring MPLS Pseudowire for Subscribers	
Chapter 27	Configuring MPLS Pseudowire Subscriber Logical Interfaces	307
	Pseudowire Subscriber Logical Interfaces Overview	307
	Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview	310
	Configuring a Pseudowire Subscriber Logical Interface	313
	Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router	315
	Configuring a Pseudowire Subscriber Logical Interface Device	315
	Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface	318
	Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces	318
	Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces	319
	Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface	321

Chapter 28	Configuring Hierarchical CoS Pseudowire Subscriber Interfaces	323
	Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview	323
	CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces	324
Chapter 29	Configuring CoS Two-Level Hierarchical Scheduling	327
	CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces	327
	Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces	329
Chapter 30	Configuring CoS Three-Level Hierarchical Scheduling	331
	CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces	331
	Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface	332
	Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set	332
	Three-Level Scheduling Hierarchy Combined Deployment Scenario	333
	Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface)	334
	Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set)	336
Part 6	Configuring Wi-Fi Access Gateway	
Chapter 31	Configuring Wi-Fi Access Gateway	341
	Wi-Fi Access Gateway Overview	342
	Wi-Fi Access Gateway Deployment Model Overview	344
	Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway	345
	Dynamic VLAN-Tagged Subscribers	345
	Untagged Subscribers	345
	Wi-Fi Access Gateway Configuration Overview	346
	Configuring a Psuedowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway	347
	Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation . . .	349
	Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways	350
	Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways	354
Part 7	Troubleshooting	
Chapter 32	Configuring PPP Log Files	359
	Configuring the Number and Size of PPP Service Log Files	359
	Configuring Access to the PPP Service Log File	360
	Configuring the Severity Level to Filter Which PPP Service Messages Are Logged	360

	Configuring a Regular Expression for PPP Service Messages to Be Logged	361
Chapter 33	Configuring PPP Trace Flags and Operations	363
	Tracing PPP Service Operations for Subscriber Access	363
	Configuring the PPP Service Trace Log Filename	364
	Configuring the PPP Service Tracing Flags	364
	Configuring Subscriber Filtering for PPP Service Trace Operations	365
Chapter 34	Configuring L2TP Log Files	367
	Configuring the Number and Size of L2TP Log Files	367
	Configuring Access to the L2TP Log File	368
	Configuring a Regular Expression for L2TP Messages to Be Logged	368
	Configuring the Severity Level to Filter Which L2TP Messages Are Logged	368
Chapter 35	Configuring L2TP Trace Flags and Operations	371
	Tracing L2TP Operations for Subscriber Access	371
	Configuring the L2TP Trace Log Filename	372
	Configuring the L2TP Tracing Flags	372
	Configuring Subscriber Filtering for L2TP Trace Operations	373
Chapter 36	Contacting Juniper Networks Technical Support	375
	Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support	375
Part 8	Configuration Statements and Operational Commands	
Chapter 37	Configuration Statements	381
	aaa-access-profile (L2TP LNS)	387
	aaa-context (AAA Options)	388
	aaa-options (Access Profile)	389
	aaa-options (PPP Profile)	390
	access (Dynamic Access Routes)	392
	access-internal (Dynamic Access-Internal Routes)	393
	access-line-information (L2TP)	394
	access-profile (AAA Options)	395
	address (L2TP Destination)	396
	address (L2TP Tunnel Destination)	397
	address (LNS Local Gateway)	398
	address (Tunnel Profile Remote Gateway)	398
	address (Tunnel Profile Source Gateway)	399
	address-change-immediate-update	399
	aggregated-inline-services-options (Aggregated Inline Services)	400
	allow-snooped-clients	401
	always-write-option-82	402
	anchor-point (Pseudowire Subscriber Interfaces)	403
	assignment-id-format (L2TP LAC)	404
	authentication (Static and Dynamic PPP)	405
	avp (L2TP Tunnel Switching)	406
	bandwidth (Inline Services)	407
	bandwidth (Tunnel Services)	408
	bearer-type (L2TP Tunnel Switching)	409

bfd	410
calling-number (L2TP Tunnel Switching)	411
challenge-length (Static and Dynamic PPP)	412
chap	413
chap (Dynamic PPP)	414
chap (L2TP)	414
cisco-nas-port-info (L2TP Tunnel Switching)	415
client	416
delimiter (Access Profile)	418
destination (L2TP)	419
destination-equal-load-balancing (L2TP LAC)	420
destruct-timeout (L2TP)	421
detection-time	422
device-count (Pseudowire Subscriber Interfaces)	423
dhcp-local-server	424
dhcp-relay	431
dhcpv6 (DHCP Local Server)	443
dhcpv6 (DHCP Relay Agent)	447
dial-options	453
dial-options (Dynamic Profiles)	454
disable-calling-number-avp (L2TP LAC)	455
disable-failover-protocol (L2TP)	456
drain	457
dual-stack-group (DHCP Local Server)	458
dual-stack-group (DHCP Relay Agent)	460
duplicate-clients (DHCPv6 Local Server and Relay Agent)	462
duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)	464
dynamic-profile (L2TP)	465
dynamic-profile (PPP)	466
enable-ipv6-services-for-lac (L2TP)	467
enable-snmp-tunnel-statistics (L2TP)	468
encapsulation (Logical Interface)	469
enforce-strict-scale-limit-license (Subscriber Management)	473
equals (Dynamic Profile)	473
failover-resync	474
failover-within-preference (L2TP LAC)	475
failure-action	476
flexible-vlan-tagging	477
forward-snooped-clients (DHCP Local Server)	478
forward-snooped-clients (DHCP Relay Agent)	479
fpc (MX Series 3D Universal Edge Routers)	480
gateway-name (LNS Local Gateway)	481
gateway-name (Tunnel Profile Remote Gateway)	482
gateway-name (Tunnel Profile Source Gateway)	482
gres-route-flush-delay (Subscriber Management)	483
group (DHCP Local Server)	484
group (DHCP Relay Agent)	487
group-profile (Group Profile)	491
hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)	493

holddown-interval	494
hello-interval (L2TP)	495
identification (Tunnel Profile)	495
idle-timeout (Access)	496
idle-timeout (L2TP)	497
ignore-magic-number-mismatch (Access Group Profile)	498
ignore-magic-number-mismatch (Dynamic Profiles)	500
initiate-ncp (Dynamic and Static PPP)	502
inline-services (FPC Level)	503
inline-services (PIC level)	504
input-hierarchical-policer	505
interface (Dynamic Routing Instances)	505
interface (L2TP Service Interfaces)	506
interface-id	507
interfaces (Static and Dynamic Subscribers)	508
ip-address-change-notify	513
ip-reassembly	514
ip-reassembly (L2TP)	515
ip-reassembly-rules (Service Set)	516
ipcp-suggest-dns-option	517
keepalive	518
keepalives	519
keepalives (Dynamic Profiles)	520
l2tp	521
l2tp (Profile)	524
l2tp-access-profile	525
l2tp-maximum-session (Service Interfaces)	526
layer2-liveness-detection (Receive)	527
layer2-liveness-detection (Send)	528
lcp-renegotiation	530
liveness-detection	531
local-gateway (L2TP LNS)	532
lockout-timeout (L2TP Destination Lockout)	533
logical-system (Tunnel Profile)	534
mac	534
mac-address (Dynamic Access-Internal Routes)	535
match-direction (IP Reassembly Rule)	536
maximum-sessions (L2TP)	537
maximum-sessions-per-tunnel	538
max-sessions (Tunnel Profile)	539
medium (Tunnel Profile)	539
method	540
metric (Dynamic Access-Internal Routes)	542
minimum-interval	543
minimum-receive-interval	544
minimum-retransmission-timeout (L2TP Tunnel)	545
mtu	546
multiplier	550
name (L2TP Destination)	551

name (L2TP Tunnel Destination)	552
no-adaptation	553
nas-port-method (L2TP LAC)	554
nas-port-method (Tunnel Profile)	554
next-hop (Dynamic Access-Internal Routes)	555
next-hop-service	556
no-allow-snooped-clients	557
no-gratuitous-arp-request	558
no-snoop (DHCP Local Server and Relay Agent)	559
no-vlan-id-validate	560
on-demand-ip-address	561
overrides (DHCP Relay Agent)	562
overrides (Enhanced Subscriber Management)	564
override-result-code (L2TP Profile)	565
pap	566
pap (Dynamic PPP)	567
pap (L2TP)	567
parse-direction (Access Profile)	568
pic (M Series and T Series Routers)	569
pool (L2TP Service Interfaces)	570
pp0 (Dynamic PPPoE)	571
ppp (Group Profile)	573
ppp-options	574
ppp-options (Dynamic PPP)	576
ppp-options (L2TP)	578
preference (Subscriber Management)	579
preference (Tunnel Profile)	580
primary-interface (Aggregated Inline Services)	581
profile (Access)	582
proxy-mode	587
ps0 (Pseudowire Subscriber Interfaces)	588
pseudowire-service (Pseudowire Subscriber Interfaces)	589
qualified-next-hop (Subscriber Management)	590
reject-unauthorized-ipv6cp	591
relay-option-82	592
remote-gateway (Tunnel Profile)	593
report-ingress-shaping-rate (Dynamic CoS Interfaces)	594
request services l2tp destination unlock	595
retransmission-count-established (L2TP)	596
retransmission-count-not-established (L2TP)	597
route (Access)	598
route (Access Internal)	599
route-suppression (DHCP Local Server and Relay Agent)	600
routing-instance (Tunnel Profile)	601
routing-instance (L2TP Destination)	601
routing-instance (L2TP Tunnel Destination)	602
routing-instances (Dynamic Profiles)	603
routing-options (Dynamic Profiles)	605
rule (IP Reassembly)	606

rx-connect-speed-when-equal (L2TP LAC)	607
rx-window-size (L2TP)	608
secondary-interface (Aggregated Inline Services)	609
secret (Tunnel Profile)	610
service-device-pool (L2TP)	610
service-device-pools (L2TP Service Interfaces)	611
service-interface (L2TP Processing)	612
service-profile (L2TP)	613
service-rate-limiter (Access)	615
session-mode	616
session-options	617
sessions-limit-group (L2TP)	618
sessions-limit-group (L2TP Client Profile)	619
shared-secret	619
soft-gre	620
source-gateway (Tunnel Profile)	621
stacked-vlan-tagging	622
statistics (Access Profile)	622
strip-user-name (Access Profile)	623
subscriber-context (AAA Options)	624
subscriber-management (Subscriber Management)	625
tag (Access)	626
threshold (detection-time)	627
threshold (transmit-interval)	628
tos-reflect (L2TP)	629
trace (DHCP Relay Agent)	630
traceoptions (Services L2TP)	631
traceoptions (Protocols PPP Service)	635
traceoptions (Subscriber Management)	638
transmit-interval	639
tunnel (L2TP)	640
tunnel (Tunnel Profile)	641
tunnel-group	642
tunnel-profile (L2TP Tunnel Switching)	643
tunnel-profile (Tunnel Profile)	644
tunnel-switch-profile (L2TP Tunnel Switching, Application)	645
tunnel-switch-profile (L2TP Tunnel Switching, Definition)	646
tx-address-change (L2TP LAC)	647
tx-connect-speed-method (L2TP LAC)	648
type (Tunnel Profile)	651
unit (Dynamic PPPoE)	652
untagged	653
user-group-profile	654
version (BFD)	655
weighted-load-balancing (L2TP LAC)	656
vlan-id (Dynamic Profiles)	657
vlan-tagging	658
vlan-tagging (Dynamic)	659
vlan-tags	660

Chapter 38	Operational Commands	661
	clear services l2tp destination	663
	clear services l2tp destination lockout	665
	clear services l2tp session	667
	clear services l2tp session statistics	670
	clear services l2tp tunnel	672
	clear services l2tp tunnel statistics	674
	request interface (revert switchover) (Aggregated Inline Service Interfaces) . .	676
	restart	678
	show bfd subscriber session	688
	show dynamic-profile session	693
	show interfaces ps0 (Pseudowire Subscriber Interfaces)	698
	show interfaces redundancy	703
	show ppp interface	706
	show ppp statistics	715
	show ppp summary	722
	show services inline ip-reassembly statistics	724
	show services l2tp client	730
	show services l2tp destination	732
	show services l2tp destination lockout	736
	show services l2tp session	738
	show services l2tp session-limit-group	747
	show services l2tp summary	749
	show services l2tp tunnel	755
	show services l2tp tunnel-group	761
	show services l2tp tunnel-switch destination	763
	show services l2tp tunnel-switch session	767
	show services l2tp tunnel-switch summary	772
	show services l2tp tunnel-switch tunnel	774
	show services soft-gre tunnel	779
	show subscribers	782
	show subscribers summary	812
	show system subscriber-management statistics	818
	show system subscriber-management summary	824
	test services l2tp tunnel	827

List of Figures

Part 1	Broadband Subscriber Access Network Overview	
Chapter 1	Broadband Subscriber Access Network Overview	3
	Figure 1: Subscriber Access Network Example	4
	Figure 2: Choosing an MSAN Type	5
	Figure 3: Basic Control Flow of Pseudowire Autosensing	7
	Figure 4: Layer 2 Services for Pseudowire Service on Service Logical Interface	10
Part 2	Configuring the DHCP Access Network	
Chapter 2	Configuring Services for DHCP Subscribers	23
	Figure 5: Subscriber Access Operation Flow	25
Chapter 9	Configuring High Availability in the DHCP Access Network	85
	Figure 6: Layer 2 Liveness Detection Send Behavior Flow	99
	Figure 7: Layer 2 Liveness Detection Receive Behavior Flow	100
Part 4	Configuring the L2TP Access Network	
Chapter 17	L2TP and Subscriber Access Overview	147
	Figure 8: Typical L2TP Topology	147
	Figure 9: Protocol Stacking for L2TP Subscribers in Pass-Through Mode	148
Chapter 18	Configuring L2TP Tunneling and Switching for Subscribers	153
	Figure 10: L2TP Tunnel Switching Network Topology	154
	Figure 11: L2TP Tunnel Switching for Incoming Calls	155
Chapter 21	Configuring L2TP LAC Tunneling for Subscribers	183
	Figure 12: Destination and Tunnel Selection Process with Failover Between Preference Levels	188
Chapter 22	Configuring Use of Subscriber Access Line and Connect Speed Information	213
	Figure 13: Sample L2TP Network Topology	214
Part 5	Configuring MPLS Pseudowire for Subscribers	
Chapter 27	Configuring MPLS Pseudowire Subscriber Logical Interfaces	307
	Figure 14: MPLS Access Network with Subscriber Management Support	308
	Figure 15: Pseudowire Subscriber Interface Protocol Stack	309
	Figure 16: Pseudowire Subscriber Logical Interface Stacking over Redundant Logical Tunnel Interface	312

Chapter 29	Configuring CoS Two-Level Hierarchical Scheduling	327
	Figure 17: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration	328
Chapter 30	Configuring CoS Three-Level Hierarchical Scheduling	331
	Figure 18: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface	332
	Figure 19: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set	333
	Figure 20: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces—Deployment Scenario	333
Part 6	Configuring Wi-Fi Access Gateway	
Chapter 31	Configuring Wi-Fi Access Gateway	341
	Figure 21: MX Series Router Deployed as a WAG	342
	Figure 22: MX Series as Wi-Fi Access Gateway Deployment Model	344

List of Tables

	About the Documentation	xxiii
	Table 1: Notice Icons	xxv
	Table 2: Text and Syntax Conventions	xxvi
Part 1	Broadband Subscriber Access Network Overview	
Chapter 1	Broadband Subscriber Access Network Overview	3
	Table 3: Ethernet MSAN Aggregation Methods	15
Part 2	Configuring the DHCP Access Network	
Chapter 5	Providing Security in the DHCP Network	41
	Table 4: Actions for DHCP Local Server Snooped Packets	43
	Table 5: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled	49
	Table 6: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled	49
	Table 7: Actions for Snooped BOOTREPLY Packets	49
Part 3	Configuring the PPP Access Network	
Chapter 14	Configuring PPP Network Control Protocol Negotiation	129
	Table 8: PPP NCP Negotiation Mode Behavior for Dynamic and Static Subscribers	130
Part 4	Configuring the L2TP Access Network	
Chapter 17	L2TP and Subscriber Access Overview	147
	Table 9: L2TP Terms	150
Chapter 18	Configuring L2TP Tunneling and Switching for Subscribers	153
	Table 10: Cause of CDN Message	156
	Table 11: Cause of StopCCN Message	156
	Table 12: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands	156
	Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary . . .	158
Chapter 21	Configuring L2TP LAC Tunneling for Subscribers	183
	Table 14: Scenario 1, Chassis Limit	200
	Table 15: Scenario 2, Tunnel Limit	200
	Table 16: Scenario 3, Tunnel Group Limit	201
	Table 17: Scenario 4, Session-Limit Group Limit	202

	Table 18: Scenario 5, Individual Client Limit	203
Chapter 22	Configuring Use of Subscriber Access Line and Connect Speed Information	213
	Table 19: L2TP AVPs That Provide Subscriber Access Line Information	214
	Table 20: Methods for Determining Connect Speeds by Junos OS Release.	222
	Table 21: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Release 17.2 and Higher)	224
	Table 22: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Releases 15.1, 16.1, 16.2, 17.1)	226
	Table 23: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Releases 13.3, 14.1, 14.2)	226
Chapter 23	Configuring L2TP LNS Inline Service Interfaces	239
	Table 24: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example	263
Chapter 26	Monitoring and Managing L2TP for Subscriber Access	299
	Table 25: SNMP Counters for L2TP Statistics	303
Part 5	Configuring MPLS Pseudowire for Subscribers	
Chapter 29	Configuring CoS Two-Level Hierarchical Scheduling	327
	Table 26: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes	327
Chapter 30	Configuring CoS Three-Level Hierarchical Scheduling	331
	Table 27: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels	331
Part 8	Configuration Statements and Operational Commands	
Chapter 38	Operational Commands	661
	Table 28: show bfd subscriber session Output Fields	688
	Table 29: show interfaces ps0 Output Fields	698
	Table 30: show interfaces redundancy Output Fields	703
	Table 31: show ppp interface Output Fields	706
	Table 32: show ppp statistics Output Fields	715
	Table 33: show ppp summary Output Fields	722
	Table 34: show services inline ip-reassembly statistics Output Fields	725
	Table 35: show services l2tp client Output Fields	730
	Table 36: show services l2tp destination Output Fields	733
	Table 37: show services l2tp destination lockout Output Fields	736
	Table 38: show services l2tp session Output Fields	739
	Table 39: show services l2tp session-limit-group Output Fields	747
	Table 40: show services l2tp summary Output Fields	749
	Table 41: show services l2tp tunnel Output Fields	756
	Table 42: show services l2tp tunnel-group Output Fields	761
	Table 43: show services l2tp tunnel-switch destination Output Fields	763
	Table 44: show services l2tp tunnel-switch session Output Fields	767
	Table 45: show services l2tp tunnel-switch summary Output Fields	772

Table 46: show services l2tp tunnel-switch tunnel Output Fields	774
Table 47: show services soft-gre tunnel Output Fields	780
Table 48: show subscribers Output Fields	786
Table 49: show subscribers summary Output Fields	813
Table 50: show system subscriber-management statistics Output Fields	818
Table 51: show system subscriber-management summary Output Fields	824
Table 52: test services l2tp tunnel Output Fields	827

About the Documentation

- Documentation and Release Notes on page xxiii
- Supported Platforms on page xxiii
- Using the Examples in This Manual on page xxiii
- Documentation Conventions on page xxv
- Documentation Feedback on page xxvii
- Requesting Technical Support on page xxvii

Documentation and Release Notes

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

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 <https://www.juniper.net/books>.

Supported Platforms

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

- MX 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 [CLI Explorer](#).

Documentation Conventions

Table 1 on page xxv 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 xxvi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

GUI Conventions

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

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

Documentation Feedback

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

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

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 <https://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 <https://www.juniper.net/support/requesting-support.html>.

PART 1

Broadband Subscriber Access Network Overview

- [Broadband Subscriber Access Network Overview on page 3](#)

CHAPTER 1

Broadband Subscriber Access Network Overview

- [Subscriber Access Network Overview on page 3](#)
- [Multiservice Access Node Overview on page 4](#)
- [LDP Pseudowire Autosensing Overview on page 5](#)
- [Layer 2 Services on Pseudowire Service Interface Overview on page 8](#)
- [Ethernet MSAN Aggregation Options on page 15](#)
- [Broadband Access Service Delivery Options on page 17](#)
- [Broadband Delivery and FTTx on page 19](#)

Subscriber Access Network Overview

A subscriber access environment can include various components, including subscriber access technologies and authentication protocols.

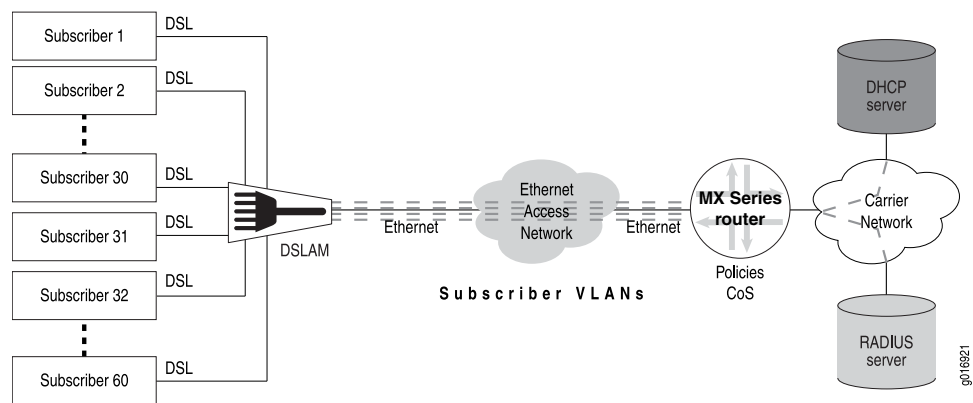
The subscriber access technologies include:

- Dynamic Host Configuration Protocol (DHCP) server
 - Local DHCP server
 - External DHCP server
- Point-to-Point Protocol (PPP)

The subscriber authentication protocols include the RADIUS server.

[Figure 1 on page 4](#) shows an example of a basic subscriber access network.

Figure 1: Subscriber Access Network Example



Related Documentation • [Subscriber Management Overview](#)

Multiservice Access Node Overview

A *multiservice access node* is a broader term that refers to a group of commonly used aggregation devices. These devices include digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections. Modern MSANs often support all of these connections, as well as providing connections for additional circuits such as plain old telephone service (referred to as POTS) or Digital Signal 1 (DS1 or T1).

The defining function of a multiservice access node is to aggregate traffic from multiple subscribers. At the physical level, the MSAN also converts traffic from the *last mile technology* (for example, ADSL) to Ethernet for delivery to subscribers.

You can broadly categorize MSANs into three types based on how they forward traffic in the network:

- **Layer-2 MSAN**—This type of MSAN is essentially a Layer 2 switch (though typically not a fully functioning switch) with some relevant enhancements. These MSANs use Ethernet (or ATM) switching to forward traffic. The MSAN forwards all subscriber traffic upstream to an edge router that acts as the centralized control point and prevents direct subscriber-to-subscriber communication. Ethernet Link Aggregation (LAG) provides the resiliency in this type of network.

Layer 2 DSLAMs cannot interpret IGMP, so they cannot selectively replicate IPTV channels.


- **Layer-3 aware MSAN**—This IP-aware MSAN can interpret and respond to IGMP requests by locally replicating a multicast stream and forwarding the stream to any subscriber requesting it. Layer 3 awareness is important when supporting IPTV traffic to perform channel changes (sometimes referred to as *channel zaps*). Static IP-aware MSANs always receive all multicast television channels. They do not have the ability to request that specific channels be forwarded to the DSLAM. Dynamic IP-aware DSLAMs, however, can inform the network to begin (or discontinue) sending individual

channels to the DSLAM. Configuring IGMP proxy or IGMP snooping on the DSLAM accomplishes this function.

- **Layer-3 MSAN**—These MSANs use IP routing functionality rather than Layer 2 technologies to forward traffic. The advantage of this forwarding method is the ability to support multiple upstream links going to different upstream routers and improving network resiliency. However, to accomplish this level of resiliency, you must assign a separate IP subnetwork to each MSAN, adding a level of complexity that can be more difficult to maintain or manage.

In choosing a MSAN type, refer to [Figure 2 on page 5](#):

Figure 2: Choosing an MSAN Type



9017267

Related Documentation

- [Ethernet MSAN Aggregation Options on page 15](#)

LDP Pseudowire Autosensing Overview

A pseudowire is a virtual link that is used to transport a Layer 2 service across an MPLS edge or access network. In a typical broadband edge or business edge network, one end of a pseudowire is terminated as a Layer 2 circuit on an access node, and the other end is terminated as a Layer 2 circuit on a service node that serves as either an aggregation node or an MPLS core network. Traditionally, both endpoints are provisioned manually through configuration. LDP pseudowire autosensing introduces a new provisioning model that allows pseudowire endpoints to be automatically provisioned and deprovisioned on service nodes based on LDP signaling messages. This model can facilitate the provisioning of pseudowires on a large scale. An access node uses LDP to signal both pseudowire identity and attributes to a service node. The identity is authenticated by a RADIUS server, and then used together with the attributes signaled by LDP and the attributes passed down by the RADIUS server to create the pseudowire endpoint configuration, including the Layer 2 circuit.

- [Pseudowire Ingress Termination Background on page 6](#)
- [Pseudowire Autosensing Approach on page 7](#)
- [Sample Configuration on page 8](#)

Pseudowire Ingress Termination Background

In a seamless MPLS-enabled broadband access or business edge network, Ethernet pseudowires are commonly used as virtual interfaces to connect access nodes to service nodes. Each pseudowire carries the bidirectional traffic of one or multiple broadband subscribers or business edge customers between an access node and a service node pair. The establishment of the pseudowire is usually initiated by the access node, based on either static configuration or dynamic detection of a new broadband subscriber or business edge customer arriving on a client-facing port on the access node.

Ideally, the access node should create one pseudowire per client port, where all subscribers or customers hosted by the port are mapped to the pseudowire. The alternative is where there is one pseudowire per client port (S-VLAN), and all subscribers or customers sharing a common S-VLAN on the port are mapped to the pseudowire. In either case, the pseudowire is signaled in the raw mode.

The S-VLAN, if not used to delimit service on the service node or combined with C-VLAN to distinguish subscribers or customers, will be stripped off before the traffic is encapsulated in pseudowire payload and transported to the service node. Individual subscribers or customers may be distinguished by C-VLAN, or a Layer 2 header such as DHCP and PPP, which will be carried in pseudowire payload to the service node. On the service node, the pseudowire is terminated. Individual subscribers or customers are then demultiplexed and modeled as broadband subscriber interfaces, business edge interfaces (for example, PPPoE), Ethernet interfaces, or IP interfaces. Ethernet and IP interfaces may be further attached to service instances, such as VPLS and Layer 3 VPN instances.

In Junos OS, pseudowire ingress termination on service nodes is supported through the use of pseudowire service physical and logical interfaces. This approach is considered as superior in scalability to the old logical tunnel interface based approach, due to its capability of multiplexing and demultiplexing subscribers or customers over a single pseudowire. For each pseudowire, a pseudowire service physical interface is created on a selected Packet Forwarding Engine, which is called an anchor Packet Forwarding Engine. On top of this pseudowire service physical interface, a ps.0 logical interface (transport logical interface) is created, and a Layer 2 circuit or Layer 2 VPN is created to host the ps.0 logical interface as an attachment interface.

The Layer 2 circuit or Layer 2 VPN enables pseudowire signaling towards the access node, and the ps.0 logical interface serves the role of customer edge facing interface for the pseudowire. Further, one or multiple ps.n logical interfaces (also known as service logical interfaces, where $n > 0$) may be created on the pseudowire service physical interface to model individual subscriber/customer flows as logical interfaces. These interfaces can then be attached to desired broadband and business edge services or Layer 2 or Layer 3 VPN instances.




NOTE: Note that the purpose of the anchor Packet Forwarding Engine is to designate the Packet Forwarding Engine to process the bidirectional traffic of the pseudowire, including encapsulation, decapsulation, VLAN mux or demux, QoS, policing, shaping, and many more.

For Junos OS Release 16.2 and earlier, the creation and deletion of the pseudowire service physical interfaces, pseudowire service logical interfaces, Layer 2 circuits, and Layer 2 VPNs for pseudowire ingress termination rely on static configuration. This is not considered as the best option from the perspective of scalability, efficiency, and flexibility, especially in a network where each service node may potentially host a large number of pseudowires. The objective is to help service providers come out of static configuration in provisioning and deprovisioning pseudowire ingress termination on service nodes.

Pseudowire Autosensing Approach

In the pseudowire autosensing approach, a service node uses the LDP label mapping message received from an access node as a trigger to dynamically generate configuration for a pseudowire service physical interface, a pseudowire service logical interface, a Layer 2 circuit. Likewise, it uses the LDP label withdraw message received from the access node and LDP session down event as triggers to remove the generated configuration. In pseudowire autosensing, it is assumed that access nodes are the initiators of pseudowire signaling, and service nodes are the targets. In a network where a service may be hosted by multiple service nodes for redundancy or load balancing, this also provides access nodes with a select-and-connect model for service establishment. The basic control flow of pseudowire autosensing is shown in [Figure 3 on page 7](#)

Figure 3: Basic Control Flow of Pseudowire Autosensing



The basic control flow procedure of pseudowire autosensing is as follows:

1. Customer premises equipment (CPE) comes online and sends an Ethernet frame with C-VLAN to the optical line terminator (OLT). OLT adds S-VLAN to the frame and sends the frame to the access node. The access node checks with the RADIUS server to authorize the VLANs.
2. The RADIUS server sends an access accept to the access node. The access node creates a Layer 2 circuit and signals a pseudowire to the service node through an LDP label mapping message.

3. The service node accepts the label mapping message, and sends an access request with pseudowire information to the RADIUS server for authorization and for selection of a pseudowire service physical interface or a logical interface.
4. The RADIUS server sends an access accept to the service node with a service string specifying the selected pseudowire service physical interface or logical interface. The service node creates a Layer 2 circuit configuration, the pseudowire information, and the pseudowire service physical interface or logical interface. The service node signals the pseudowire towards the access node through an LDP label mapping message. The pseudowire comes up bidirectionally.

Sample Configuration

The following configuration explicitly marks the Layer 2 circuit as generated by autosensing. The pseudowire service physical interface and pseudowire service logical interface configuration are optional, depending on whether they preexist.

```
Router O [edit]
          protocols {
            Layer 2 circuit {
              neighbor 192.0.2.2 {
                interface ps0.0 {
                  virtual-circuit-id 100;
                  control-word;
                  mtu 9100;
                  auto-sensed;
                }
              }
            }
          }
```

- Related Documentation**
- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
 - [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Layer 2 Services on Pseudowire Service Interface Overview

The pseudowire service logical interface supports the transport logical interface (psn.0) on the MPLS access side and service logical interfaces (psn.1 to psn.n) on the MPLS core side of the subscriber management network.

The pseudowire service on service logical interfaces psn.1 to psn.n are configured as Layer 2 interfaces in the bridge domain or in a virtual private LAN service (VPLS) instance. There is Layer 2 circuit or the Layer 2 VPN across MPLS access between an Ethernet aggregation device and a service edge device with the pseudowire service on transport logical interface ps n.0 as the terminating interface of the Layer 2 circuit or the Layer 2 VPN at the service edge device.

Junos OS supports the pseudowire service on service logical interfaces psn.1 to psn.n in the bridge domain or VPLS instance, which receives traffic egressing from the pseudowire

service on the transport logical interface at the service edge device. It also enables Layer 2 ingress features such as MAC learning, VLAN manipulations, and destination MAC look up on the pseudowire service on service logical interfaces.

When the traffic is in reverse direction, the destination MAC enters the Layer 2 domain at the service edge device, which is learned as the source MAC on the pseudowire service on service logical interfaces. Junos OS supports the Ethernet VPLS, Ethernet bridge, VLAN VPLS, and VLAN bridge encapsulation next hop on the pseudowire service on service logical interfaces to exit Layer 2 traffic. The Layer 2 output features such as VLAN manipulations and others are enabled on the pseudowire service on service logical interfaces. The traffic sent out of the pseudowire service on service logical interfaces enter the pseudowire service on transport logical interfaces which is the Layer 2 circuit interface between Ethernet aggregation and service edge devices across the MPLS access domain.



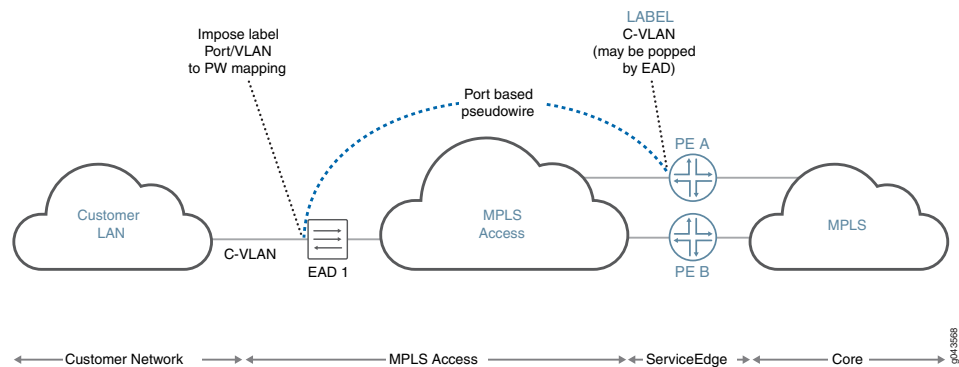
NOTE: For Junos OS Release 16.2 and earlier, Layer 2 encapsulations or features could not be configured on pseudowire service on service logical interfaces.

Traffic from Customer LAN to MPLS

VPLS-x and VPLS-y instances are configured on the MPLS core side of the service edge device (PE A). A Layer 2 circuit or Layer 2 VPN is configured between the Ethernet aggregation device (EAD1) and the service edge device. ps0.0 (transport logical interface) is the local interface in the Layer 2 circuit or the Layer 2 VPN at PE A. Junos OS supports pseudowire service on service logical interface ps0.x (x>0) in VPLS instance VPLS-x (VLAN ID in VPLS-x = m) and pseudowire service on service logical interface ps0.y (y>0) in VPLS instance VPLS-y (VLAN ID in VPLS-y = n).

In [Figure 4 on page 10](#), when the traffic comes from EAD1 to PE A (on either Layer 2 circuit or Layer 2 VPN) with any VLAN ID, the traffic will exit through ps0.0. Based on the VLAN ID in the traffic the pseudowire service on service logical interface is selected. For example, if VLAN ID is m, then the traffic will enter ps0.x and if VLAN ID is n, then the traffic will enter ps0.y.

Figure 4: Layer 2 Services for Pseudowire Service on Service Logical Interface



When traffic enters pseudowire service on the service logical interface ps0.n, where $n > 0$, the following steps are performed.

1. The source MAC learning should occur on the Layer 2 pseudowire service on the service logical interface. The source Packet Forwarding Engine for this MAC is the Packet Forwarding Engine of the logical tunnel interface on which the pseudowire service is anchored in a VPLS instance or bridge domain in the PE A device.
2. The destination MAC lookup is done at the entry side as an input bridge family feature list of pseudowire services on service logical interfaces.
 - If destination MAC lookup is successful, then the traffic is sent as unicast; otherwise, the destination MAC, broadcast MAC, and multicast MAC are flooded.
 - If destination MAC lookup fails for the traffic coming on a pseudowire service on a service logical interface, the **mlp query** command is sent to the Routing Engine and the other Packet Forwarding Engine in bridge domain or VPLS instance.
3. If a new MAC is learned on a pseudowire service on a service logical interface, then the **mlp add** command is sent to the Routing Engine and the other Packet Forwarding Engine in bridge domain or VPLS instance.

Traffic from Service Edge to Customer LAN

When traffic enters the VPLS instance or bridge domain at the service edge device and if the destination MAC in the traffic is learned on a pseudowire service on a service logical interface, then the token associated with that pseudowire service logical interface is set at the entry side. The traffic is then sent to the Packet Forwarding Engine on which the logical tunnel interface of the pseudowire service physical interface is anchored through a fabric. When this token is launched, it supports VLAN VPLS, VLAN bridge, Ethernet VPLS, and Ethernet bridge encapsulations. The encapsulation next hop points to the egress logical interface feature list of the pseudowire service on the service logical interface to execute all the Layer 2 output features and send the packet to the entry side of the pseudowire service on transport logical interface ps0.0.

If the MAC query reaches the Packet Forwarding Engine on which the pseudowire service is anchored, then the Packet Forwarding Engine sends the response only when the MAC learned on the pseudowire service on the service logical interface is present. The Layer 2 token associated with the pseudowire service on the service logical interface seen after destination MAC lookup for the MAC learned on the pseudowire service on service logical interface should point to the next hop associated with the access side of the pseudowire service on service the logical interface.

The pseudowire service on the transport logical interface is the local interface ps0.0 of the Layer 2 circuit or Layer 2 VPN between the service edge and the Ethernet aggregation devices. Traffic is sent to the Ethernet aggregation device though the Layer 2 circuit or Layer 2 VPN across the MPLS access domain.

If the destination MAC traffic coming from the entry and exit side of the service edge device is unknown or multicast or broadcast, the traffic needs to be flooded. This requires an customer edge device flood next hop to include the pseudowire service on service logical interface, which acts as an access logical interface for the VPLS instance or bridge domain.

Pseudowire Service Interfaces

The following features are supported on pseudowire service interfaces:

- A pseudowire service interface is hosted over a logical tunnel interface (lt-x/y/z). The traffic from a transport pseudowire service on a logical interface to a subscriber pseudowire service on a logical interface is based on the available VLAN ID.
- Transfer of traffic from a subscriber pseudowire service on a logical interface to a transport pseudowire service on a logical interface is based on the channelID through an available loopback IP address.
- Pseudowire service on service logical interfaces are supported on the virtual routing and forwarding (VRF) routing instance.

Sample Configuration

This sample configuration shows a pseudowire service on a transport logical interface on a Layer 2 circuit and a pseudowire service on service logical interfaces in a bridge domain and a VPLS instance in a service edge device:

**Pseudowire service on
a service logical
interface in bridge
domain on router 0**

```
[edit]
interfaces {
  ps0 {
    unit 0 {
      encapsulation Ethernet-ccc;
    }
    unit 1 {
      encapsulation VLAN-bridge;
      VLAN-id 1;
    }
    unit 2 {
      encapsulation VLAN-bridge;
      VLAN-id 2;
    }
  }
}
```

```
    }
  }
  ge-0/0/0 {
    unit 1 {
      encapsulation VLAN-bridge;
      VLAN-id 1;
    }
    unit 2 {
      encapsulation VLAN-bridge;
      VLAN-id 2;
    }
  }
  ge-2/0/6 {
    unit 0 {
      family inet {
        address 10.11.2.1/24;
      }
      family iso;
      family mpls;
    }
  }
}
protocols {
  mpls {
    label-switched-path to_192.0.2.2 {
      to 192.0.2.2;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.3.3;
    }
  }
  Layer 2 circuit {
    neighbor 192.0.2.2 {
      interface ps0.0 {
        virtual-circuit-id 100;
      }
    }
  }
}
bridge-domains {
  bd1 {
    domain-type bridge;
    VLAN-id 1;
    interface ps0.1;
    interface ge-0/0/0.1;
  }
  bd2 {
    domain-type bridge;
    VLAN-id 2;
    interface ps0.2;
    interface ge-0/0/0.2;
  }
}
```


Pseudowire service on
a service logical
interface in a VPLS
instance on router 0

```
[edit]
interfaces {
  ps0 {
    unit 0 {
      encapsulation Ethernet-ccc;
    }
    unit 1 {
      encapsulation VLAN-vpls;
      VLAN-id 1;
      family vpls;
    }
    unit 2 {
      encapsulation VLAN-vpls;
      VLAN-id 2;
      family vpls;
    }
  }
  ge-0/0/0 {
    unit 1 {
      encapsulation VLAN-vpls;
      VLAN-id 1;
      family vpls;
    }
    unit 2 {
      encapsulation VLAN-vpls;
      VLAN-id 2;
      family vpls;
    }
  }
  ge-2/0/6 {
    unit 0 {
      family inet {
        address 10.11.2.1/24;
      }
      family iso;
      family mpls;
    }
  }
}
protocols {
  mpls {
    label-switched-path to_192.0.2.2 {
      to 192.0.2.2;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.3.3;
    }
  }
  Layer 2 circuit {
    neighbor 192.0.2.2 {
      interface ps0.0 {
        virtual-circuit-id 100;
      }
    }
  }
}
```

```
    }  
  }  
}  
routing-instances {  
  vpls-1 {  
    instance-type vpls;  
    VLAN-id 1;  
    interface ps0.1;  
    interface ge-0/0/0.1;  
  }  
  vpls-2 {  
    instance-type vpls;  
    VLAN-id 2;  
    interface ps0.2;  
    interface ge-0/0/0.2;  
  }  
}
```

**Pseudowire service on
a service logical
interface in a Layer 2
circuit on router 0**

```
[edit]  
interfaces {  
  ps0 {  
    unit 0 {  
      encapsulation Ethernet-ccc;  
    }  
    unit 1 {  
      encapsulation VLAN-ccc;  
      VLAN-id 1;  
    }  
    unit 2 {  
      encapsulation VLAN-ccc;  
      VLAN-id 2;  
    }  
  }  
  ge-0/0/0 {  
    unit 1 {  
      encapsulation VLAN-vpls;  
      VLAN-id 1;  
      family vpls;  
    }  
    unit 2 {  
      encapsulation VLAN-vpls;  
      VLAN-id 2;  
      family vpls;  
    }  
  }  
  ge-2/0/6 {  
    unit 0 {  
      family inet {  
        address 10.11.2.1/24;  
      }  
      family iso;  
      family mpls;  
    }  
  }  
}
```

```
protocols {
  mpls {
    label-switched-path to_192.0.2.2 {
      to 192.0.2.2;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.3.3;
    }
  }
  Layer 2 circuit {
    neighbor 192.0.2.2 {
      interface ps0.0 {
        virtual-circuit-id 100;
      }
    }
    neighbor 10.10.10.10 {
      interface ps0.1 {
        virtual-circuit-id 1;
      }
    }
    neighbor 10.11.11.11 {
      interface ps0.2 {
        virtual-circuit-id 2;
      }
    }
  }
}
```

- Related Documentation
- [Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface on page 321](#)
 - [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
 - [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Ethernet MSAN Aggregation Options

Each MSAN can connect directly to an edge router (broadband services router or video services router), or an intermediate device (for example, an Ethernet switch) can aggregate MSAN traffic before being sent to the services router. [Table 3 on page 15](#) lists the possible MSAN aggregation methods and under what conditions they are used.

Table 3: Ethernet MSAN Aggregation Methods

Method	When Used
Direct connection	Each MSAN connects directly to the broadband services router and optional video services router.

Table 3: Ethernet MSAN Aggregation Methods (*continued*)

Method	When Used
Ethernet aggregation switch connection	Each MSAN connects directly to an intermediate Ethernet switch. The switch, in turn, connects to the broadband services router or optional video services router.
Ethernet ring aggregation connection	Each MSAN connects to a ring topology of MSANs. The head-end MSAN (the device closest to the upstream edge router) connects to the broadband services router.

You can use different aggregation methods in different portions of the network. You can also create multiple layers of traffic aggregation within the network. For example, an MSAN can connect to a central office terminal (COT), which, in turn, connects to an Ethernet aggregation switch, or you can create multiple levels of Ethernet aggregation switches prior to connecting to the edge router.

Direct Connection

In the direct connection method, each MSAN has a point-to-point connection to the broadband services router. If an intermediate central office exists, traffic from multiple MSANs can be combined onto a single connection using wave-division multiplexing (WDM). You can also connect the MSAN to a video services router. However, this connection method requires that you use a Layer 3 MSAN that has the ability to determine which link to use when forwarding traffic.

When using the direct connection method, keep the following in mind:

- We recommend this approach when possible to simplify network management.
- Because multiple MSANs are used to connect to the services router, and Layer 3 MSANs generally require a higher equipment cost, this method is rarely used in a multiedge subscriber management model.
- Direct connection is typically used when most MSAN links are utilized less than 33 percent and there is little value in combining traffic from multiple MSANs.

Ethernet Aggregation Switch Connection

An Ethernet aggregation switch aggregates traffic from multiple downstream MSANs into a single connection to the services router (broadband services router or optional video services router).

When using the Ethernet aggregation switch connection method, keep the following in mind:

- Ethernet aggregation is typically used when most MSAN links are utilized over 33 percent or to aggregate traffic from lower speed MSANs (for example, 1 Gbps) to a higher speed connection to the services router (for example, 10 Gbps).
- You can use an MX Series router as an Ethernet aggregation switch. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Ring Aggregation Connection

In a ring topology, the remote MSAN that connects to subscribers is called the remote terminal (RT). This device can be located in the outside plant (OSP) or in a remote central office (CO). Traffic traverses the ring until it reaches the central office terminal (COT) at the head-end of the ring. The COT then connects directly to the services router (broadband services router or video services router).



NOTE: The RT and COT must support the same ring resiliency protocol.

You can use an MX Series router in an Ethernet ring aggregation topology. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Related Documentation

- [Multiservice Access Node Overview on page 4](#)

Broadband Access Service Delivery Options

Four primary delivery options exist today for delivering broadband network service. These options include the following:

- [Digital Subscriber Line on page 17](#)
- [Active Ethernet on page 18](#)
- [Passive Optical Networking on page 18](#)
- [Hybrid Fiber Coaxial on page 18](#)

Digital Subscriber Line

Digital subscriber line (DSL) is the most widely deployed broadband technology worldwide. This delivery option uses existing telephone lines to send broadband information on a different frequency than is used for the existing voice service. Many generations of DSL are used for residential service, including Very High Speed Digital Subscriber Line 2 (VDSL2) and versions of Asymmetric Digital Subscriber Line (ADSL, ADSL2, and ADSL2+). These variations of DSL primarily offer asymmetric residential broadband service where different upstream and downstream speeds are implemented. (VDSL2 also supports symmetric operation.) Other DSL variations, like High bit rate Digital Subscriber Line (HDSL) and Symmetric Digital Subscriber Line (SDSL), provide symmetric speeds and are typically used in business applications.

The head-end to a DSL system is the Digital Subscriber Line Access Multiplexer (DSLAM). The demarcation device at the customer premise is a DSL modem. DSL service models are defined by the Broadband Forum (formerly called the DSL Forum).

Active Ethernet

Active Ethernet uses traditional Ethernet technology to deliver broadband service across a fiber-optic network. Active Ethernet does not provide a separate channel for existing voice service, so VoIP (or TDM-to-VoIP) equipment is required. In addition, sending full-speed (10 or 100 Mbps) Ethernet requires significant power, necessitating distribution to Ethernet switches and optical repeaters located in cabinets outside of the central office. Due to these restrictions, early Active Ethernet deployments typically appear in densely populated areas.

Passive Optical Networking

Passive Optical Networking (PON), like Active Ethernet, uses fiber-optic cable to deliver services to the premises. This delivery option provides higher speeds than DSL but lower speeds than Active Ethernet. Though PON provides higher speed to each subscriber, it requires a higher investment in cable and connectivity.

A key advantage of PON is that it does not require any powered equipment outside of the central office. Each fiber leaving the central office is split using a non-powered optical splitter. The split fiber then follows a point-to-point connection to each subscriber.

PON technologies fall into three general categories:

- ATMPON (APON), Broadband PON (BPON), and Gigabit-capable PON (GPON)—PON standards that use the following different delivery options:
 - APON—The first passive optical network standard is primarily used for business applications.
 - BPON—Based on APON, BPON adds wave division multiplexing (WDM), dynamic and higher upstream bandwidth allocation, and a standard management interface to enable mixed-vendor networks.
 - GPON—The most recent PON adaptation, GPON is based on BPON but supports higher rates, enhanced security, and a choice of which Layer 2 protocol to use (ATM, Generic Equipment Model [GEM], or Ethernet).
- Ethernet PON (EPON)—Provides capabilities similar to GPON, BPON, and APON, but uses Ethernet standards. These standards are defined by the IEEE. Gigabit Ethernet PON (GEAPON) is the highest speed version.
- Wave Division Multiplexing PON (WDM-PON)—A nonstandard PON which, as the name implies, provides a separate wavelength to each subscriber.

The head-end to a PON system is an Optical Line Terminator (OLT). The demarcation device at the customer premises is an Optical Network Terminator (ONT). The ONT provides subscriber-side ports for connecting Ethernet (RJ-45), telephone wires (RJ-11) or coaxial cable (F-connector).

Hybrid Fiber Coaxial

Multi-System Operators (MSOs; also known as *cable TV operators*) offer broadband service through their hybrid fiber-coaxial (HFC) network. The HFC network combines

optical fiber and coaxial cable to deliver service directly to the customer. Services leave the central office (CO) using a fiber-optic cable. The service is then converted outside of the CO to a coaxial cable *tree* using a series of optical nodes and, where necessary, through a trunk radio frequency (RF) amplifier. The coaxial cables then connect to multiple subscribers. The demarcation device is a cable modem or set-top box, which talks to a Cable Modem Termination System (CMTS) at the MSO *head-end* or master facility that receives television signals for processing and distribution. Broadband traffic is carried using the Data Over Cable Service Interface Specification (DOCSIS) standard defined by CableLabs and many contributing companies.

Related Documentation • [Broadband Delivery and FTTx on page 19](#)

Broadband Delivery and FTTx

Many implementations use existing copper cabling to deliver signal to the premises, but fiber-optic cable connectivity is making its way closer to the subscriber. Most networks use a combination of both copper and fiber-optic cabling. The term *fiber to the x* (FTTx) describes how far into the network fiber-optic cabling runs before a switch to copper cabling takes place. Both PON and Active Ethernet can use fiber-optic portion of the network, while xDSL is typically used on the copper portion. This means that a single fiber-optic strand may support multiple copper-based subscribers.

Increasing the use of fiber in the network increases cost but it also increases network access speed to each subscriber.

The following terms are used to describe the termination point of fiber-optic cable in a network:

- Fiber to the Premises (FTTP), Fiber to the Home (FTTH), Fiber to the Business (FTTB)—Fiber extends all the way to the subscriber. PON is most common for residential access, although Active Ethernet can be efficiently used in dense areas such as apartment complexes. Active Ethernet is more common for delivering services to businesses.
- Fiber to the Curb (FTTC)—Fiber extends most of the way (typically, 500 feet/150 meters or less) to the subscriber. Existing copper is used for the remaining distance to the subscriber.
- Fiber to the Node/Neighborhood (FTTN)—Fiber extends to within a few thousand feet of the subscriber and converted to xDSL for the remaining distance to the subscriber.
- Fiber to the Exchange (FTTE)—A typical central office-based xDSL implementation in which fiber is used to deliver traffic to the central office and xDSL is used on the existing local loop.

Related Documentation • [Broadband Access Service Delivery Options on page 17](#)

PART 2

Configuring the DHCP Access Network

- [Configuring Services for DHCP Subscribers on page 23](#)
- [Applying RADIUS Route Attributes to Subscribers or to Access Networks on page 31](#)
- [Suppressing DHCP Access, Access-Internal, and Destination Routes on page 37](#)
- [Providing Security in the DHCP Network on page 41](#)
- [Distinguishing Between Duplicate DHCPv4 Subscribers on the Same Subnet on page 61](#)
- [Distinguishing Between Duplicate DHCPv6 Subscribers on page 67](#)
- [Using the DHCP Relay Agent to Selectively Process DHCP Client Traffic on page 71](#)
- [Configuring High Availability in the DHCP Access Network on page 85](#)
- [Monitoring and Managing DHCP for Subscriber Access on page 111](#)

CHAPTER 2

Configuring Services for DHCP Subscribers

- [DHCP and Subscriber Management Overview on page 23](#)
- [Subscriber Access Operation Flow Using DHCP Relay on page 24](#)
- [Defining Various Levels of Services for DHCP Subscribers on page 25](#)
- [Example: Configuring a Tiered Service Profile for Subscriber Access on page 26](#)

DHCP and Subscriber Management Overview

You use DHCP in broadband access networks to provide IP address configuration and service provisioning. DHCP, historically a popular protocol in LANs, works well with Ethernet connectivity and is becoming increasingly popular in broadband networks as a simple, scalable solution for assigning IP addresses to subscriber home PCs, set-top boxes (STBs), and other devices.

Junos OS subscriber management supports the following DHCP allocation models:

- DHCP Local Server
- DHCP Relay
- DHCP Relay Proxy

DHCP uses address assignment pools from which to allocate subscriber addresses. Address-assignment pools support both dynamic and static address assignment:

- Dynamic address assignment—A subscriber is automatically assigned an address from the address-assignment pool.
- Static address assignment—Addresses are reserved and always used by a particular subscriber.



NOTE: Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

Extended DHCP Local Server and Subscriber Management Overview

You can enable the services router to function as an extended DHCP local server. As an extended DHCP local server the services router, and not an external DHCP server, provides

an IP address and other configuration information in response to a client request. The extended DHCP local server supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Extended DHCP Relay and Subscriber Management Overview

You can configure extended DHCP relay options on the router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You can use DHCP relay in carrier edge applications such as video and IPTV to obtain configuration parameters, including an IP address, for your subscribers. The extended DHCP relay agent supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

DHCP Relay Proxy and Subscriber Management Overview

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits. Except for the ability to add DHCP relay agent options and the gateway address (giaddr) to DHCP packets, DHCP relay is transparent to DHCP clients and DHCP servers, and simply forwards messages between DHCP clients and servers. When you configure DHCP relay to operate in proxy mode, the relay is no longer transparent. In proxy mode, DHCP relay conceals DHCP server details from DHCP clients, which interact with a DHCP relay in proxy mode as though it is the DHCP server. For DHCP servers there is no change, because proxy mode has no effect on how the DHCP server interacts with the DHCP relay.


- Related Documentation**
- *Extended DHCP Local Server Overview*
 - *Extended DHCP Relay Agent Overview*
 - *DHCP Relay Proxy Overview*
 - *Address-Assignment Pools Overview*

Subscriber Access Operation Flow Using DHCP Relay

The subscriber management feature requires that a subscriber (for example, a DHCP client) send a discover message to the router interface to initialize dynamic configuration of that interface.

[Figure 5 on page 25](#) shows the flow of operations that occurs when the router is using DHCP relay to enable access for a subscriber.

Figure 5: Subscriber Access Operation Flow



The following general sequence occurs during access configuration for a DHCP client:

1. The client issues a DHCP discover message.
2. The router issues an authorization request to the RADIUS server.
3. The RADIUS server issues an authorization response to the router.
4. The router passes the DHCP discover message through to the DHCP server.
5. The DHCP server issues an IP address for the client.
6. The router DHCP component sends an acknowledgement back to the client.

The subscriber now has access to the network and the authorized service.

- Related Documentation**
- *Subscriber Management Overview*
 - *Configuring Subscriber Access*

Defining Various Levels of Services for DHCP Subscribers

This topic discusses how to create dynamic profiles to define various levels of service for DHCP clients.

Before you configure dynamic profiles for client services:

1. Create a basic dynamic profile.
See Configuring a Basic Dynamic Profile.
2. Configure a dynamic profile that enables DHCP clients access to the network.
See Configuring Dynamic DHCP Client Access to a Multicast Network



NOTE: You can create a basic dynamic profile that contains both access configuration and some level of basic service.

3. Ensure that the router is configured to enable communication between the client and the RADIUS server.

See [Specifying the Authentication and Accounting Methods for Subscriber Access](#).

4. Configure all RADIUS values that you want the profiles to use when validating DHCP clients.

See [Configuring RADIUS Server Parameters for Subscriber Access](#)

To configure an initial client access dynamic profile:

1. Access the desired service profile.

```
user@host# set dynamic-profiles basic-service-profile
```

2. (Optional) Define any IGMP protocols values as described for creating a basic access profile to combine a basic service with access in a profile.

See [Configuring Dynamic DHCP Client Access to a Multicast Network](#).

3. (Optional) Specify any filters for the interface.

See [Dynamically Attaching Statically Created Filters for Any Interface Type](#), [Dynamically Attaching Statically Created Filters for a Specific Interface Family Type](#), or [Dynamically Attaching Filters Using RADIUS Variables](#).

4. Define any CoS values for the service level you want this profile to configure on the interface.

Related Documentation

- [Configuring a Basic Dynamic Profile](#)
- [Dynamic Profiles Overview](#)

Example: Configuring a Tiered Service Profile for Subscriber Access

This example shows how to configure a tiered service profile for subscribers.

The profile contains three services:

- Gold—Subscribers that pay for this service are allocated 10M bandwidth for data, voice, and video services.
- Silver—Subscribers that pay for this service are allocated 5M bandwidth for data, voice, and video services.

- Bronze—Subscribers that pay for this service are allocated 1M bandwidth for the data service only.

Each subscriber is allocated a VLAN that is created statically. Subscribers log in using DHCP and authenticate using RADIUS. The subscribers can migrate from one service to another when they change subscriptions.

To configure a profile for a tiered service:

1. Configure the VLAN interfaces associated with each subscriber. Enable hierarchical scheduling for the interface.

```
interfaces {
  ge-2/0/0 {
    description subscribers;
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      vlan-tags outer 100 inner 100;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 127.0.0.2;
      }
    }
    unit 2 {
      family inet {
        vlan-tags outer 101 inner 101;
        unnumbered-address lo0.0 preferred-source-address 127.0.0.2;
      }
    }
    unit 3 {
      vlan-tags outer 102 inner 102;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 127.0.0.2;
      }
    }
  }
}
```

2. Configure the static CoS parameters.

In this example, each offering (video, voice, and data) is assigned a queue, and each service (Gold, Silver, and Bronze) is assigned a scheduler.

```
class-of-service {
  forwarding-classes {
    queue 0 data;
    queue 1 voice;
    queue 2 video;
  }
  scheduler-maps {
    bronze_service_smap {
      forwarding-class data scheduler data_sch;
    }
    silver_service_smap {
      forwarding-class data scheduler data_sch;
    }
  }
}
```

```
        forwarding-class voice scheduler silver_voice_sch;
        forwarding-class video scheduler silver_video_sch;
    }
    gold_service_smap {
        forwarding-class data scheduler data_sch;
        forwarding-class voice scheduler gold_voice_sch;
        forwarding-class video scheduler gold_video_sch;
    }
}
schedulers {
    data_sch {
        transmit-rate percent 20;
        buffer-size remainder;
        priority low;
    }
    silver_voice_sch {
        transmit-rate percent 30;
        buffer-size remainder;
        priority high;
    }
    silver_video_sch {
        transmit-rate percent 30;
        buffer-size remainder;
        priority medium;
    }
    gold_voice_sch {
        transmit-rate percent 40;
        buffer-size remainder;
        priority high;
    }
    gold_video_sch {
        transmit-rate percent 40;
        buffer-size remainder;
        priority medium;
    }
}
}
```

3. Configure the dynamic profile for the service.

The scheduler maps configured for each service are referenced in the dynamic profile.

```
dynamic-profiles {
    subscriber_profile {
        interfaces {
            "$junos-interface-ifd-name" {
                unit "$junos-underlying-interface-unit" {
                    family inet;
                }
            }
        }
    }
    class-of-service {
        traffic-control-profiles {
            subscriber_tcp {
                scheduler-map $smap;
                shaping-rate $shaping-rate;
            }
        }
    }
}
```



```

        guaranteed-rate $guaranteed-rate;
        delay-buffer-rate $delay-buffer-rate;
    }
}
interfaces {
    "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile subscriber_tcp;
        }
    }
}
}
}

```

4. Configure access for the subscribers.

The DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You use DHCP relay to obtain configuration parameters, including an IP address, for subscribers. In this example, one DHCP server, address 198.51.100.1, can be used by subscribers.

The DHCP relay configuration is attached to an active server group named `service_provider_group`.

The subscribers are grouped together within the `subscriber_group`, and identifies characteristics such as authentication, username info, and the associated interfaces for the group members. In this example, it also identifies the active server group and the dynamic interface that is used by the subscribers in the group.

```

forwarding-options {
    dhcp-relay {
        server-group {
            service_provider_group {
                198.51.100.1;
            }
        }
        group subscriber_group {
            active-server-group service_provider_group;
            dynamic-profile subscriber_profile;
            interface ge-2/0/0.1;
            interface ge-2/0/0.2;
            interface ge-2/0/0.3;
        }
    }
}
}

```

Related Documentation

- For more information about configuring CoS for subscriber access, see *CoS for Subscriber Access Overview*

CHAPTER 3

Applying RADIUS Route Attributes to Subscribers or to Access Networks

- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32](#)
- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

Access and Access-Internal Routes for Subscriber Management

DHCP and PPP on the router use both access routes and access-internal routes to represent either the subscriber or the networks behind the attached router. An access route represents a network behind an attached router, and is set to a preference of 13. An access-internal route is a /32 route that represents a directly attached subscriber, and is set to a preference of 12.

Access routes typically are used to apply the values of the RADIUS Framed-Route attribute [22] for IPv4 routes and the Framed-IPv6-Route attribute [99] for IPv6 routes. A framed route consists of a prefix that represents a public network behind the CPE, a next-hop gateway, and optional route attributes consisting of a combination of metric, preference, and tag. The only mandatory component of the framed route is the prefix. The next-hop gateway can be specified explicitly in the framed route, as 0.0.0.0, ::0, or the subscriber's fixed address assigned by the Framed-IP-Address (8) or Framed-IPv6-Prefix (97) attribute (common practice for business subscribers). Alternatively, the absence of the gateway address implies address 0.0.0.0. The address 0.0.0.0 or ::0, whether implicit or explicitly configured, resolves to the subscriber's assigned address (host route). Consequently, the convention is that the next-hop gateway is the subscriber's IP address.

You can configure a dynamic profile to use predefined variables to dynamically configure access routes using the values specified in the RADIUS attribute. To configure access routes include the **access** stanza at the **[edit dynamic-profiles profile-name routing-options]** hierarchy level.

Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes in the dynamic profile configuration. If the RADIUS Framed-Route attribute (22) or Framed-IPv6-Route attribute [99] does not specify the next-hop gateway—as is common—the variable representing the next-hop, \$junos-framed-route-nexthop or

`$junos-framed-route-ipv6-next-hop`, automatically resolves to the subscriber's IP address. If you configure the **access-internal** statement in the dynamic profile, it is ignored.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)
- [RADIUS IETF Attributes Supported by the AAA Service Framework](#)

Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management

You can dynamically configure access-internal routes. In releases earlier than Junos OS 15.1, this configuration is optional, but can be included so that the values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

Starting in Junos OS Release 15.1R1, we no longer recommend that you always include the **access-internal** stanza in the dynamic-profile when the **access** stanza is present for framed route support. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute (22) or Framed-IPv6-Route attribute [99].

DHCP subscriber interfaces require the qualified-next-hop to identify the interface and the MAC address.

To dynamically configure access-internal routes:

1. Specify that you want to configure the access-internal route.

```
user@host# edit dynamic-profiles profile-name routing-options
```

2. Configure the IP address and the qualified next-hop address as variables.

```
[edit dynamic-profiles profile-name routing-options]
user@host# edit access-internal route $junos-subscriber-ip-address qualified-next-hop
$junos-interface-name
```



NOTE: The variable used for **qualified-next-hop** is `$junos-interface-name`.

3. Configure the MAC address for the qualified next-hop as a variable.

```
[edit dynamic-profiles profile-name routing-options access-internal route
$junos-subscriber-ip-address qualified-next-hop $junos-underlying-interface]
user@host# set mac-address $junos-subscriber-mac-address
```

Release History Table

Release	Description
15.1R1	Starting in Junos OS Release 15.1R1, we no longer recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

Related Documentation

- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)
- [Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 111](#)

Configuring Dynamic Access Routes for Subscriber Management

You can dynamically configure access routes for DHCP and PPP subscribers based on the values specified in the following RADIUS attributes:

- For IPv4 access routes, use the variable, **\$junos-framed-route-ip-address-prefix**. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, **\$junos-framed-route-ipv6-address-prefix**. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].

To dynamically configure access routes:

1. Configure the route prefix for the access route as a variable.

For IPv4:

```
[edit dynamic-profiles profile-name routing-options]
user@host# edit access route $junos-framed-route-ip-address-prefix
```

For IPv6:

```
[edit dynamic-profiles profile-name routing-options]
user@host# edit access route $junos-framed-route-ipv6-address-prefix
```

2. Configure the next-hop address as a variable.

For IPv4:

```
[edit dynamic-profiles profile-name routing-options access route
"$junos-framed-route-ip-address-prefix"]
user@host# set next-hop $junos-framed-route-nexthop
```

For IPv6:

```
[edit dynamic-profiles profile-name routing-options access route
"$junos-framed-route-ipv6-address-prefix"]
user@host# set next-hop $junos-framed-route-ipv6-nexthop
```

3. Configure the metric as a variable.

For IPv4:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set metric $junos-framed-route-cost
```

For IPv6:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set metric $junos-framed-route-ipv6-cost
```

4. Configure the preference as a variable.

For IPv4:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set preference $junos-framed-route-distance
```

For IPv6:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set preference $junos-framed-route-ipv6-distance
```

5. Configure the tag as a variable.

IPv4:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set tag $junos-framed-route-tag
```

IPv6:

```
[edit dynamic-profiles profile-name routing-options access route  
"$junos-framed-route-ip-address-prefix"]  
user@host# set tag $junos-framed-route-ipv6-tag
```

Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If the RADIUS Framed-Route attribute (22) or Framed-IPv6-Route attribute [99] does not specify the next-hop gateway—as is common—the variable representing the next-hop, `$junos-framed-route-nexthop`, is automatically resolved. If you configure the **access-internal** statement in the dynamic profile, it is ignored.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

**Related
Documentation**

- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32](#)
- [Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 123](#)
- [Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 111](#)
- *RADIUS IETF Attributes Supported by the AAA Service Framework*

CHAPTER 4

Suppressing DHCP Access, Access-Internal, and Destination Routes

- [Suppressing DHCP Access, Access-Internal, and Destination Routes on page 37](#)
- [Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 38](#)

Suppressing DHCP Access, Access-Internal, and Destination Routes

During the DHCP client binding operation, the DHCP process adds route information for the DHCP sessions by default. The DHCP process adds access-internal and destination routes for DHCPv4 sessions, and access-internal and access routes for DHCPv6 sessions. In some scenarios, you might want to override the default behavior and prevent DHCP from automatically installing the route information. For example, DHCP relay installs destination (host) routes by default—this action is required in certain configurations to enable address renewals from the DHCP server to work properly. However, the default installation of destination routes might cause a conflict when you configure DHCP relay with static subscriber interfaces. To avoid such configuration conflicts you can override the default behavior and prevent DHCP relay from installing the routes.



NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.

You can configure both DHCP local server and DHCP relay agent to override the default route installation behavior, and you can specify the override for both DHCPv4 and DHCPv6 sessions. You can override the route installation globally or for named interface groups. For DHCPv4 you can override the installation of destination routes only or access-internal routes (the access-internal option prevents installation of both destination and access-internal routes). For DHCPv6 you can specify access routes, access-internal routes, or both.

Related Documentation

- [Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 38](#)
- [Extended DHCP Local Server Overview](#)

- [DHCPv6 Local Server Overview](#)
- [Extended DHCP Relay Agent Overview](#)
- [DHCPv6 Relay Agent Overview](#)

Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default

You can configure both DHCP local server and DHCP relay agent to override the default installation of access, access-internal, and destination routes. For DHCPv4 you can override the installation of destination routes only or access-internal routes (the access-internal option prevents installation of both destination and access-internal routes). For DHCPv6 you can specify access routes, access-internal routes, or both. You can configure the override globally or for named interface groups.



NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.



NOTE: The `no-arp` statement is deprecated and the function is replaced by the `route-suppression` statement.

To configure route suppression and prevent DHCP from installing specific types of routes:

- For DHCP local server route suppression (for example, a global configuration):

```
[edit system services dhcp-local-server]
user@host# set route-suppression access-internal
```

- For DHCP relay (for example, a group-specific configuration):

```
[edit forwarding-options dhcp-relay group southeast]
user@host# set route-suppression destination
```

- For DHCPv6 local server (for example, a group-specific configuration):

```
[edit system services dhcp-local-server group southern3]
user@host# set dhcpv6 route-suppression access access-internal
```

- For DHCPv6 relay (for example, a global configuration):

```
[edit forwarding-options dhcp-relay]
user@host# set dhcpv6 route-suppression access
```

Related Documentation

- [Suppressing DHCP Access, Access-Internal, and Destination Routes on page 37](#)
- [Extended DHCP Local Server Overview](#)
- [DHCPv6 Local Server Overview](#)
- [Extended DHCP Relay Agent Overview](#)

- *DHCPv6 Relay Agent Overview*

CHAPTER 5

Providing Security in the DHCP Network

- [DHCP Snooping Support on page 41](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 43](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48](#)
- [Disabling DHCP Snooping Filters on page 51](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 52](#)
- [Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent on page 54](#)
- [Preventing DHCP Spoofing on page 58](#)

DHCP Snooping Support

DHCP snooping provides DHCP security by identifying incoming DHCP packets. In the default DHCP snooping configuration, all traffic is snooped. You can optionally use the **forward-snooped-clients** statement to evaluate the snooped traffic and to determine if the traffic is forwarded or dropped, based on whether or not the interface is configured as part of a group.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. The router discards snooped packets by default if there is no subscriber associated with the packet. To enable normal processing of snooped packets, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for a specific routing instance for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets. A renew request may be unicast directly to the DHCP server. This is a BOOTPREQUEST packet and is snooped.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.
- You can also disable snooping filters. In the preceding configurations, all DHCP traffic is forwarded to the slower routing plane of the routing instance before it is either forwarded or dropped. Disabling snooping filters causes DHCP traffic that can be forwarded directly from the faster hardware control plane to bypass the routing control plane.

Related Documentation

- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 43](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48](#)
- [Disabling DHCP Snooping Filters on page 51](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 52](#)
- [Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent on page 54](#)

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

You can configure how DHCP local server handles DHCP snooped packets. Depending on the configuration, DHCP local server either forwards or drops the snooped packets it receives.

Table 4 on page 43 indicates the action the router takes for DHCP local server snooped packets.



NOTE: Configured interfaces are those interfaces that have been configured with the `group` statement in the `[edit system services dhcp-local-server]` hierarchy. Non-configured interfaces are those that are in the logical system/routing instance but have not been configured by the `group` statement.

Table 4: Actions for DHCP Local Server Snooped Packets

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	forwarded	forwarded
configured-interfaces	forwarded	dropped
non-configured-interfaces	dropped	forwarded

To configure DHCP snooped packet forwarding for DHCP local server:

1. Specify that you want to configure DHCP local server.

```
[edit]
user@host# edit system services dhcp-local-server
```

2. Enable DHCP snooped packet forwarding for DHCP local server.

```
[edit system services dhcp-local-server]
user@host# edit forward-snooped-clients
```

3. Specify the interfaces that are supported for snooped packet forwarding.

```
[edit system services dhcp-local-server forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP local server to forward DHCP snooped packets on only configured interfaces:

```
[edit]
```

```
system {
  services {
    dhcp-local-server {
      forward-snooped-clients configured-interfaces;
    }
  }
}
```

Related Documentation

- [DHCP Snooping Support on page 41](#)

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes the first procedure, in which you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration.

The second procedure, which applies only to DHCPv4 relay agent, is described in [“Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent” on page 48](#), and configures the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

You can enable or disable DHCP globally for DHCP relay, for a group of interfaces, or for a specific interface in a group.

By default, DHCP snooping is disabled for DHCP relay. To enable or disable DHCP snooping support globally:

1. Specify that you want to configure DHCP relay agent.
 - For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```
 - For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```
2. Specify that you want to override the default configuration.
 - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```
 - For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```
3. Enable or disable DHCP snooping support.
 - To enable DHCP snooping:
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay overrides]
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-allow-snooped-clients
```

For example, to enable global DHCP snooping support :

```
forwarding-options {
  dhcp-relay {
    overrides {
      allow-snooped-clients;
    }
  }
}
```

To enable or disable DHCP snooping support for a group of interfaces:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit group group-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group group-name
```

3. Specify that you want to override the default configuration.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]  
user@host# edit overrides
```

4. Enable or disable DHCP snooping support.

- To enable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]  
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]  
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]  
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]  
user@host# set no-allow-snooped-clients
```

For example, to enable DHCP snooping support on all interfaces in group **boston**:

```
forwarding-options {  
  dhcp-relay {  
    group boston {  
      overrides {  
        allow-snooped-clients;  
      }  
    }  
  }  
}
```

To enable or disable DHCP snooping support on a specific interface:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]  
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]  
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group containing the interface.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]  
user@host# edit group group-name
```

- For DHCPv6 relay agent:


```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group group-name
```
- 3. Specify the interface for which you want to configure DHCP snooping.
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit interface interface-name
```
 - For DHCPv6 relay agent:


```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
user@host# edit interface interface-name
```
- 4. Specify that you want to override the default configuration on the interface.
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay group group-name interface interface-name]
user@host# edit overrides
```
 - For DHCPv6 relay agent:


```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name]
user@host# edit overrides
```
- 5. Enable or disable DHCP snooping support.
 - To enable DHCP snooping:
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay group group-name interface interface-name
overrides]
user@host# set allow-snooped-clients
```
 - For DHCPv6 relay agent:


```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name overrides]
user@host# set allow-snooped-clients
```
 - To disable DHCP snooping:
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay group group-name interface interface-name
overrides]
user@host# set no-allow-snooped-clients
```
 - For DHCPv6 relay agent:


```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name overrides]
user@host# set no-allow-snooped-clients
```

For example, to disable DHCP snooping support on interface **ge-2/1/8.0** in group **boston**:

```
forwarding-options {
```

```

dhcp-relay {
  group boston {
    interface ge-2/1/8.0 {
      overrides {
        no-allow-snooped-clients;
      }
    }
  }
}

```

To enable DHCPv6 snooping support on interface **ge-3/2/1.1** in group **sunnyvale**:

```

forwarding-options {
  dhcp-relay {
    dhcpv6 {
      group sunnyvale {
        interface ge-3/2/1.1 {
          overrides {
            allow-snooped-clients;
          }
        }
      }
    }
  }
}

```

Related Documentation

- [DHCP Snooping Support on page 41](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 52](#)
- [Overriding the Default DHCP Relay Configuration Settings](#)

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

You can configure how DHCP relay agent handles DHCP snooped packets. Depending on the configuration, DHCP relay agent either forwards or drops the snooped packets it receives.

DHCP relay uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes how you use the **forward-snooped-clients** statement to manage whether DHCP relay agent forwards or drops snooped packets, depending on the type of interface on which the packets are snooped. In the other part of the DHCP relay agent snooping configuration, which is described in “[Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent](#)” on page 44, you enable or disable the DHCP relay snooping feature.

[Table 5 on page 49](#) shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the **allow-snooped-clients** statement. [Table 6 on page 49](#) shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the **no-allow-snooped-clients** statement.

The router or switch also uses the configuration of the DHCP relay agent forwarding support to determine how to handle snooped BOOTREPLY packets. [Table 7 on page 49](#) shows the action the router (or switch) takes for the snooped BOOTREPLY packets.



NOTE: Configured interfaces have been configured with the `group` statement in the `[edit forwarding-options dhcp-relay]` hierarchy. Non-configured interfaces are in the logical system/routing instance but have not been configured by the `group` statement.

Table 5: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
<code>forward-snooped-clients</code> not configured	snooped packets result in subscriber (DHCP client) creation	dropped
<code>all-interfaces</code>	forwarded	forwarded
<code>configured-interfaces</code>	forwarded	dropped
<code>non-configured-interfaces</code>	snooped packets result in subscriber (DHCP client) creation	forwarded

Table 6: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
<code>forward-snooped-clients</code> not configured	dropped	dropped
<code>all-interfaces</code>	dropped	forwarded
<code>configured-interfaces</code>	dropped	dropped
<code>non-configured-interfaces</code>	dropped	forwarded

Table 7: Actions for Snooped BOOTREPLY Packets

forward-snooped-clients Configuration	Action
<code>forward-snooped-clients</code> not configured	snooped BOOTREPLY packets dropped if client is not found
<code>forward-snooped-clients</code> all configurations	snooped BOOTREPLY packets forwarded if client is not found

To configure DHCP snooped packet forwarding and BOOTREPLY snooped packet forwarding for DHCP relay agent:

1. Specify that you want to configure DHCP relay agent.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

2. Enable DHCP snooped packet forwarding.

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```

3. Specify the interfaces that are supported for snooped packet forwarding.

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP relay agent to forward DHCP snooped packets on only configured interfaces:

```
[edit]
forwarding-options {
  dhcp-relay {
    forward-snooped-clients configured-interfaces;
  }
}
```

**Related
Documentation**

- [DHCP Snooping Support on page 41](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)

Disabling DHCP Snooping Filters

DHCP snooping provides DHCP security by identifying incoming DHCP packets. In the default DHCP snooping configuration, all traffic is snooped. You can optionally use the **forward-snooped-clients** statement to evaluate the snooped traffic and to determine whether the traffic is forwarded or dropped, based on whether or not the interface is configured as part of a group.

In both the default configuration and in configurations using the **forward-snooped-clients** statement, all DHCP traffic is forwarded from the hardware control plane to the routing plane of the routing instance to ensure that all DHCP packets are intercepted. In certain topologies, such as a Metropolitan Routing Ring topology, forwarding all DHCP traffic to the control plane can result in excessive traffic. The **no-snoop** configuration statement disables the snooping filter for DHCP traffic that can be directly forwarded on the hardware control plane, such as Layer 3 unicast packets with a valid route, causing those DHCP packets to bypass the slower routing plane. You can disable DHCP snooping filters starting in Junos OS Release 15.1R2.

To disable DHCP snooping filters on the DHCP local server:

1. Specify that you want to configure DHCP local server.

```
[edit]
user@host# edit system services dhcp-local-server
```

2. Disable DHCP snooping filters for DHCP local server.

```
[edit system services dhcp-local-server]
user@host# set no-snoop
```

3. Specify that you want to configure DHCPv6 local server.

```
[edit system services dhcp-local-server]
user@host# edit dhcpv6
```

4. Disable DHCP snooping filters for DHCPv6 local server.

```
[edit system services dhcp-local-server dhcpv6]
user@host# set no-snoop
```

To disable DHCP snooping filters on the DHCP relay server:

1. Specify that you want to configure DHCP relay server.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

2. Disable DHCP snooping filters for DHCP local server.

```
[edit forwarding-options dhcp-relay]
user@host# set no-snoop
```

3. Specify that you want to configure DHCPv6 relay server.

```
[edit forwarding-options dhcp-relay]
user@host# edit dhcpv6
```

4. Disable DHCP snooping filters for DHCPv6 local server.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set no-snoop
```

Release History Table

Release	Description
15.1R2	You can disable DHCP snooping filters starting in Junos OS Release 15.1R2.

Related Documentation

- [DHCP Snooping Support on page 41](#)
- [no-snoop on page 559](#)

Example: Configuring DHCP Snooping Support for DHCP Relay Agent

This example shows how to configure DHCP snooping support for DHCP relay agent.

- [Requirements on page 52](#)
- [Overview on page 52](#)
- [Configuration on page 52](#)

Requirements

- Configure DHCP relay agent. See *Extended DHCP Relay Agent Overview*.

Overview

In this example, you configure DHCP snooping support for DHCP relay agent by completing the following operations:

- Override the default DHCP snooping configuration and enable DHCP snooping support for the interfaces in group **frankfurt**.
- Configure DHCP relay agent to forward snooped packets to only configured interfaces.



NOTE: By default, DHCP snooping is disabled globally.

Configuration

Step-by-Step Procedure

To configure DHCP relay support for DHCP snooping:

1. Specify that you want to configure DHCP relay agent.


```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- Specify the named group of interfaces on which DHCP snooping is supported.

```
[edit forwarding-options dhcp-relay]
user@host# edit group frankfurt
```

- Specify the interfaces that you want to include in the group. DHCP relay agent considers these as the configured interfaces when determining whether to forward or drop traffic.

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

- Specify that you want to override the default configuration for the group.

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# edit overrides
```

- Enable DHCP snooping support for the group.

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# set allow-snooped-clients
```

- Return to the **[edit forwarding-options dhcp-relay]** hierarchy level to configure the forwarding action and specify that DHCP relay agent forward snooped packets on only configured interfaces:

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# up 2
```

- Enable DHCP snooped packet forwarding for DHCP relay agent.

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```

- Specify that snooped packets are forwarded on only configured interfaces (the interfaces in group **frankfurt**).

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set configured-interfaces
```

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group **frankfurt**.

```
[edit]
user@host# show forwarding-options
dhcp-relay {
```

```
forward-snooped-clients configured-interfaces;
group frankfurt {
  overrides {
    allow-snooped-clients;
  }
  interface fe-1/0/1.3 {
    upto fe-1/0/1.9;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related
Documentation**

- [DHCP Snooping Support on page 41](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)

Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent

Snooping support for DHCPv6 relay agent is disabled on the router by default. This example shows how to override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCPv6 snooping for a named group of interfaces and for a specific interface within a different named group.



NOTE: You can also enable DHCPv6 snooping support globally by using the `allow-snooped-clients` statement at the `[edit forwarding-options dhcp-relay dhcpv6 overrides]` hierarchy level.

-
- [Requirements on page 54](#)
 - [Overview on page 55](#)
 - [Configuration on page 55](#)
 - [Verification on page 57](#)

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers
- Junos OS Release 12.1 or later

Before you begin:

- Configure DHCPv6 relay agent.

See *DHCPv6 Relay Agent Overview*

- Configure named DHCPv6 relay agent interface groups to which you want to apply a common DHCP configuration.

See Grouping Interfaces with Common DHCP Configurations

Overview

In this example, you override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCP snooping for both of the following:

- All of the interfaces in the group named **boston**
- Interface **ge-3/2/1.1** in the group named **sunnyvale**

Configuration

To override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCPv6 snooping for a named group of interfaces and for a specific interface within a named group, perform these tasks:

- [Enabling DHCPv6 Snooping Support for a Named Group of Interfaces on page 55](#)
- [Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group on page 56](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set forwarding-options dhcp-relay dhcpv6 group boston overrides allow-snooped-clients
set forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1 overrides
allow-snooped-clients
```

Enabling DHCPv6 Snooping Support for a Named Group of Interfaces

Step-by-Step Procedure

To enable DHCPv6 snooping support for a named group of interfaces:

1. Specify that you want to configure DHCPv6 relay agent.

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```
2. Specify the named group of interfaces for which you want to enable DHCPv6 snooping.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group boston
```
3. Specify that you want to override the default DHCPv6 configuration for the interfaces in that group.

```
[edit forwarding-options dhcp-relay dhcpv6 group boston]
user@host# edit overrides
```
4. Enable DHCPv6 snooping support for all interfaces in group **boston**.

```
[edit forwarding-options dhcp-relay dhcpv6 group boston overrides]
user@host# set allow-snooped-clients
```

Results From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options dhcp-relay]
user@host# show
dhcpv6 {
  group boston {
    overrides {
      allow-snooped-clients;
    }
  }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group

Step-by-Step Procedure To enable DHCPv6 snooping support for a specific interface within a named group of interfaces:

1. Return to the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level to specify that you want to configure DHCPv6 relay agent.

```
[edit forwarding-options dhcp-relay dhcpv6 group boston overrides]
user@host# up 2
```

2. Specify the named group containing the interface.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group sunnyvale
```

3. Specify the interface in group **sunnyvale** for which you want to enable DHCPv6 snooping.

```
[edit forwarding-options dhcp-relay dhcpv6 group sunnyvale]
user@host# edit interface ge-3/2/1.1
```

4. Specify that you want to override the default DHCPv6 configuration for interface **ge-3/2/1.1** in group **sunnyvale**.

```
[edit forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1]
user@host# edit overrides
```

5. Enable DHCPv6 snooping support for interface **ge-3/2/1.1** in group **sunnyvale**.

```
[edit forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1
overrides]
```

```
user@host# set allow-snooped-clients
```

Results From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options dhcp-relay]
user@host# show
dhcpv6 {
  group boston {
    overrides {
      allow-snooped-clients;
    }
  }
  group sunnyvale {
    interface ge-3/2/1.1 {
      overrides {
        allow-snooped-clients;
      }
    }
  }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

Verification

To verify the DHCPv6 configuration in a multi-relay topology, perform this task:

- [Verifying the Address Bindings for DHCPv6 Relay Agent Clients on page 57](#)

Verifying the Address Bindings for DHCPv6 Relay Agent Clients

Purpose Verify the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Action Display detailed information about address bindings for DHCPv6 relay agent clients.

```
user@host > show dhcpv6 relay binding detail
```

```
Session Id: 13
  Client IPv6 Prefix: 2001:db8:0:8001::5/128
  Client DUID: LL0x1-00:00:65:03:01:02
  State: BOUND(DHCPV6_RELAY_STATE_BOUND)
  Lease Expires: 2011-11-21 06:14:50 PST
  Lease Expires in: 293 seconds
  Lease Start: 2011-11-21 06:09:50 PST
  Incoming Client Interface: ge-3/2/1.1
  Server Address: unknown
  Next Hop Server Facing Relay: 2001:db8::2
  Server Interface: none
  Client Id Length: 10
  Client Id: /0x00030001/0x00006503/0x0102
```

Meaning The **Server Address** field in the **show dhcpv6 relay binding detail** command output typically displays the IP address of the DHCPv6 server. In this example, the value **unknown** in the **Server Address** field indicates that this is a multi-relay topology in which the DHCPv6 relay agent is not directly adjacent to the DHCPv6 server, and does not detect the IP address of the server.

In that case, the output instead includes the **Next Hop Server Facing Relay** field, which displays the next-hop address in the direction of the DHCPv6 server.

- Related Documentation**
- [DHCPv6 Relay Agent Overview](#)
 - [DHCP Snooping Support on page 41](#)
 - [Grouping Interfaces with Common DHCP Configurations](#)
 - [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)

Preventing DHCP Spoofing

A problem that sometimes occurs with DHCP is *DHCP spoofing*. In DHCP spoofing, an untrusted client floods a network with DHCP messages. Often these attacks utilize source IP address spoofing to conceal the true source of the attack.

DHCP snooping helps prevent DHCP spoofing by copying DHCP messages to the control plane and using the information in the packets to create anti-spoofing filters. The anti-spoofing filters bind a client's MAC address to its DHCP-assigned IP address and use this information to filter spoofed DHCP messages. In a typical topology, a carrier edge router (in this function also referred to as the broadband network gateway [BNG]) connects the DHCP server and the MX Series router (or broadband services aggregator [BSA]) performing the snooping. The MX Series router connects to the client and the BNG.

To configure DHCP snooping, you include the appropriate interfaces within a DHCP group. You can configure DHCP snooping for VPLS environments and bridge domains.

- In a VPLS environment, DHCP requests are forwarded over pseudowires. You configure DHCP snooping over VPLS at the **[edit routing-instances routing-instance-name]** hierarchy level.
- In bridge domains, DHCP snooping works on a per learning bridge basis. Each learning domain must have an upstream interface configured. This interface acts as the flood port for DHCP requests coming from the client side. DHCP requests are forwarded across learning domains in a bridge domain. You configure DHCP snooping on bridge domains at the **[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]** hierarchy level.

To configure DHCP relay to prevent DHCP spoofing:

1. Access the appropriate hierarchy for either a VPLS or bridge domain configuration.

```
user@host# edit routing-instances blue
```

2. Specify that you want to configure DHCP relay.

```
[edit routing-instances blue]
user@host# edit forwarding-options dhcp-relay
```

3. Create the group and assign a name.

```
[edit routing-instances blue forwarding-options dhcp-relay]
user@host# edit group svl-10
```

4. Specify the names of one or more interfaces. DHCP will trust only the MAC addresses learned on the specified interfaces.

```
[edit routing-instances blue forwarding-options dhcp-relay group svl-10]
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```



NOTE: You can explicitly enable and disable interface support for DHCP snooped clients. See [“Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent”](#) on page 44.

Related Documentation

- *Extended DHCP Relay Agent Overview*
- For examples of DHCP snooping, see the *JUNOS MX Series Ethernet Services Routers Solutions Guide*.
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 44](#)

CHAPTER 6

Distinguishing Between Duplicate DHCPv4 Subscribers on the Same Subnet

- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)
- [Guidelines for Configuring Support for DHCPv4 Duplicate Clients on page 62](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information on page 63](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces on page 64](#)

DHCPv4 Duplicate Client In Subnet Overview

In some network environments, client IDs and hardware addresses (MAC addresses) might not be unique, resulting in duplicate clients. A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same hardware address as an existing DHCP client—the existing client and the new client cannot exist simultaneously, unless you have configured the optional duplicate client support.

By default, DHCP local server and DHCP relay agent use the subnet information to differentiate between duplicate clients. However, in some cases, this level of differentiation is not adequate. For example, when multiple subinterfaces share the same underlying loopback interface with the same preferred source address, the interfaces appear to be on the same subnet.

You can enable support for duplicate clients in a subnet by configuring DHCP to use additional information to uniquely identify clients—the additional information is either the client incoming interface or the option 82 information in the DHCP packets. Using the option 82 information provides the following important benefits:

- You can configure DHCP relay to preserve and use the remotely created option 82.
- DHCP local server can support an environment in which an aggregation device is present between the client and the DHCP server.

When configured to support duplicate clients in the subnet, DHCP uses the following information to distinguish between the duplicate clients:

- The subnet on which the client resides

- The client ID or hardware address
- The duplicate clients option you configure—either the client incoming interface or the option 82 information in the client's incoming DHCP packets



NOTE: Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are used. Other suboptions, such as Vendor-Specific (suboption 9), are ignored.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are used.

Related Documentation

- [Guidelines for Configuring Support for DHCPv4 Duplicate Clients on page 62](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces on page 64](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information on page 63](#)
- [DHCPv6 Duplicate Client DUIDs on page 67](#)

Guidelines for Configuring Support for DHCPv4 Duplicate Clients

When configuring DHCPv4 duplicate client support, consider the following guidelines:

- If you want to preserve the remotely-created option 82 information, use the **option 82** option with the **duplicate-clients-in-subnet** statement to distinguish between duplicate clients. If there is no remotely created option 82 in the incoming DHCP packets, the router locally creates the option 82 information.
- If you want to use the locally-created option-82, use the **incoming-interface** option with the **duplicate-clients-in-subnet** statement to distinguish between duplicate clients.
- Only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.
- DHCP relay agent and DHCP local server in the same routing instance must have the same the **duplicate-clients-in-subnet** configuration.
- For the Layer 3 wholesale model:
 - The wholesaler and retailer logical system/routing instances must have the same **duplicate-clients-in-subnet** statement configuration.

- For DHCP relay, the wholesaler and the retailer routing contexts must both have the **relay-option-82** statement configured with the Agent Circuit ID suboption (suboption 1) in option 82.

Related Documentation

- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces on page 64](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information on page 63](#)
- [DHCPv6 Duplicate Client DUIDs on page 67](#)

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information

Duplicate clients occur when two clients in a subnet have the same hardware address or the same client ID.

The following two procedures describe how to configure the router to use the option 82 information in the incoming packets to differentiate between duplicate clients. The first procedure describes the configuration for DHCP relay agent. The second procedure is for DHCP local server.



NOTE: Only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

To configure the DHCP relay agent to differentiate between duplicate clients based on option 82 information:

1. Specify that you want to configure DHCP relay agent.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```

2. Configure DHCP relay to insert option 82 information if there is no remotely created option 82. Use the default setting, which inserts the interface ID rather than the optional interface description.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option-82 circuit-id
```

3. Configure the router to always accept DHCP client packets that contain option 82 information.

```
[edit forwarding-options dhcp-relay]
user@host# set overrides trust-option-82
```



NOTE: The `trust-option-82` statement must always be enabled so the router can process incoming DHCP client packets that contain option 82 information when the packets have a gateway IP address (giaddr) of 0 (zero).

4. Configure DHCP relay to use the remotely created option 82 information to distinguish between duplicate clients. If there is no remotely created option 82 in the traffic, the router locally creates the option 82 information.

```
[edit forwarding-options dhcp-relay]
user@host# set duplicate-clients-in-subnet option-82
```



NOTE: Make sure that the `always-write-option-82` statement is *not* enabled, as the statement will overwrite the remotely created option 82.

To configure the DHCP local server to differentiate between duplicate clients based on the option 82 information:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Configure the duplicate client support with the `option-82` option.

```
[edit system services dhcp-local-server]
user@host# set duplicate-clients-in-subnet option-82
```

Related Documentation

- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)
- [Guidelines for Configuring Support for DHCPv4 Duplicate Clients on page 62](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces on page 64](#)

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces

Duplicate clients occur when two clients in a subnet have the same hardware address or the same client ID.

The following two procedures describe how to configure the router to use the clients' incoming interface to differentiate between duplicate clients. The first procedure describes the configuration for DHCP relay agent; the second procedure is for DHCP local server.

To configure the DHCP relay agent to differentiate between duplicate clients based on the client incoming interface:

1. Specify that you want to configure DHCP relay agent.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```

2. Configure the duplicate client support with the **incoming-interface** option.

```
[edit forwarding-options dhcp-relay]
user@host# set duplicate-clients-in-subnet incoming-interface
```

3. Configure DHCP relay to insert option 82 information if the information is not specified remotely. Use the default setting, which inserts the interface ID rather than the optional interface description.



NOTE: Only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option-82 circuit-id
```

4. Configure the router to overwrite any remotely supplied option 82 information in incoming packets.

```
[edit forwarding-options dhcp-relay]
user@host# set overrides always-write-option-82
```

5. Configure the router to always accept DHCP client packets that contain option 82 information.

```
[edit forwarding-options dhcp-relay]
user@host# set overrides trust-option-82
```



NOTE: The *trust-option-82* statement must always be enabled so the router can process incoming DHCP client packets that contain option 82 information when the packets have a gateway IP address (giaddr) of 0 (zero).

To configure the DHCP local server to differentiate between duplicate clients based on the client incoming interface:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Configure the duplicate client support with the **incoming-interface** option.

```
[edit system services dhcp-local-server]
```

```
user@host# set duplicate-clients-in-subnet incoming-interface
```

**Related
Documentation**

- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)
- [Guidelines for Configuring Support for DHCPv4 Duplicate Clients on page 62](#)
- [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information on page 63](#)

CHAPTER 7

Distinguishing Between Duplicate DHCPv6 Subscribers

- [DHCPv6 Duplicate Client DUIDs on page 67](#)
- [Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs on page 68](#)

DHCPv6 Duplicate Client DUIDs

The DHCP unique identifier (DUID) is used to identify a client for the proper application of configuration parameters. The DUID is supposed to be unique across all clients. A duplicate DHCPv6 client occurs when a client attempts to obtain a lease, and that client has the same DUID as an existing DHCPv6 client. Because the DUIDs are supposed to be unique, by default the router treats the request from the duplicate client as a renegotiation by the original client, and replaces the existing client entry with a new entry.

However, in some cases the duplicate request is legitimate, because some network equipment vendors do not guarantee the uniqueness of DUIDs. In these circumstances the router can support the duplication of the DUID by accommodating the new client without affecting the existing client.

Starting in Junos OS Release 16.1, you can enable duplicate DHCPv6 client support. When enabled, the router uses the clients' underlying (incoming) interfaces to differentiate between clients with the same DUID. The router can then create a new client entry for the duplicate client and grant it a lease. The router retains the existing client entry with the original lease.

All underlying interface types are supported. Only 1:1 VLANs are supported, because the client requests are received over different underlying interfaces. N:1 VLANs are not supported, because the client requests can be received over the same underlying interface and therefore cannot be differentiated if the DUIDs are the same.

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, you can enable duplicate DHCPv6 client support.

Related Documentation

- [Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs on page 68](#)
- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)

Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs

DHCPv6 duplicate clients occur when two clients in a subnet have the same DHCP Unique Identifier (DUID).

The following procedure describes how to configure the router to use the client's underlying (incoming) interface to differentiate between clients with duplicate DUIDs. The first part of the procedure describes the configuration for DHCPv6 relay agent and the second part configures the DHCPv6 local server.



NOTE: Duplicate client DUIDs are supported only when the clients use different underlying interfaces, as in the case of 1:1 VLANs. They are not supported when the clients share an underlying interface, as in the case of N:1 VLANs.

Before configuring duplicate client support, you must ensure the following:

- DHCPv6 relay agent is configured to insert the DHCPv6 Interface-ID option (option 18) in packets forwarded to the DHCPv6 local server.
- Option 18 specifies the interface name, not the text description of the interface.
- DHCPv6 local server must echo option 18 in the RELAY-REPLY messages returned to the DHCPv6 relay agent, as is the case for DHCPv6 local server configured on a Juniper Networks router. The relay agent uses the echoed option 18 information to find the client's interface and construct the client key.

To configure the DHCPv6 relay agent to support duplicate DUIDs:

1. Specify that you want to configure DHCPv6 relay agent.

```
[edit forwarding-options]
user@host# edit dhcp-relay dhcpv6
```

2. Configure DHCPv6 relay agent to insert DHCPv6 option 18 in the packets forwarded to the DHCPv6 local server.

```
[edit forwarding-options dhcp-relay dhcpv6]
```



```
user@host# set relay-agent-interface-id
```



NOTE: You must not include the `use-interface-description` statement because it specifies a text description of the interface.

3. Specify that the DHCPv6 relay agent uses the clients' incoming interfaces to differentiate between the duplicate DUIDs.

```
[edit forwarding-options dhcp-relay dhcpv6]  
user@host# set duplicate-clients incoming-interface
```

To configure the DHCPv6 local server to support duplicate DUIDs:

1. Specify that you want to configure DHCPv6 local server.

```
[edit system services]  
user@host# edit dhcp-local-server dhcpv6
```

2. Configure the DHCPv6 local server to support duplicate clients based on the clients' incoming interfaces.

```
[edit system services dhcp-local-server dhcpv6]  
user@host# set duplicate-clients incoming-interface
```

**Related
Documentation**

- [DHCPv6 Duplicate Client DUIDs on page 67](#)
- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)

CHAPTER 8

Using the DHCP Relay Agent to Selectively Process DHCP Client Traffic

- [DHCP Options and Selective Traffic Processing Overview on page 71](#)
- [Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 73](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 75](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 79](#)

DHCP Options and Selective Traffic Processing Overview

Subscriber management enables you to provide selective traffic processing based on information that is provided in the DHCP and DHCPv6 options string included in the traffic. Starting in Junos OS Release 15.1, the selective traffic processing feature lets you manage multivendor networks with the extended DHCP and DHCPv6 relay agent. You can enable the extended DHCP and DHCPv6 relay agent to compare option-specific strings received in DHCP client packets against a list of ASCII or hexadecimal strings that you configure on the router. The selective traffic processing feature allows you to identify traffic based on the option in the DHCP client packets, filter the traffic, and specify the action that the DHCP relay takes for the traffic. You can use DHCP options 60 and 77 and DHCPv6 options 15 and 16 to identify client traffic. You configure the action the router takes for the selected traffic, such as forwarding the traffic to a specific DHCP server, or dropping the traffic. DHCP relay agent selective traffic processing also allows you to specify a default action, which the router uses when no other action satisfies the configuration.

Using selective traffic processing is helpful in network environments where DHCP clients access services that are provided by multiple vendors and by multiple DHCP servers. For example, a DHCP client might gain Internet access from a particular DHCP server provided by one vendor, and access an IPTV service from a different DHCP server owned by a second vendor. Using the option-specific information in the DHCP client packets enables DHCP relay agent to differentiate between the two servers and to take the correct action for the subscriber.

You might also use selective processing to distinguish between services to different DHCP subscribers on the same interface. For example, a household might include two IP devices

that obtain their IP addresses from the service provider's DHCP server. The service provider might want to bind one of the devices to the incoming interface, sharing that address with other households. At the same time the service provider might want the second device to have its own filter and CoS capabilities. For this second device, the service provider can use selective processing to create a dynamic IP demux interface.

You can configure selective processing support globally or for a named group of interfaces. You can also configure the support for the extended DHCP relay agent on a per logical system and per routing instance basis.

To configure selective processing, you specify the DHCP or DHCPv6 option attribute that identifies the traffic, the match criteria used to filter the traffic, and the action to perform with the filtered traffic.

You can use the following DHCP options to selectively process client traffic:

- DHCPv4 option 60 (Vendor Class Identifier)
- DHCPv4 option 77 (User Class Identifier)
- DHCPv6 option 15 (User Class Option)
- DHCPv6 option 16 (Vendor Class Option)

You can configure exact match or partial match criteria to filter client traffic, and specify either the **ascii** option (to define a nonempty ASCII string of 1 through 255 alphanumeric characters) or the **hexadecimal** option (to define a hexadecimal string of 1 through 255 hexadecimal characters [0 through 9, a through f, and A through F]).



BEST PRACTICE: Because of the format of DHCP option 77 and DHCPv6 option 16, we recommend you configure hexadecimal matching only with these two options instead of ASCII matching.

You can configure an unlimited number of match strings. If you configure a string as both exact match (**equals**) and a partial match (**starts-with**) criteria, the exact match takes precedence. Wildcard characters are not supported in exact match or partial match strings.

Use the following match criteria to filter client traffic:

- **equals**—Your specified string is an exact match to the option string in client traffic.
- **starts-with**—Your specified string is a subset of the option string in client traffic, starting with the left-most character. For example, your configuration of the string "test" is a subset of "test123" in the client's option string, and matches the **starts-with** criteria.
- **default-action**—The option string in client traffic does not satisfy any match criteria, or no match criteria are configured.



NOTE: The default-action is optional. If the match criteria are not satisfied or not configured and there is no default-action configured, DHCP relay processes the traffic in the normal manner.

You can specify the following actions for the filtered client traffic:

- **drop**—Discard the traffic.
- **forward-only**—Forward the traffic, without creating a new subscriber session.



NOTE: When you use the forward-only action, the only configured **overrides** operation supported is the trust-option-82 option. DHCP relay agent ignores all other **overrides** options that are configured.

- **local-server-group**—Forward the traffic to the specified group of DHCP local servers that provides the requested client service. This option is not supported for DHCPv6 relay agent.
- **relay-server-group**—Forward the traffic to the specified group of DHCP servers that provides the requested client service.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, the selective traffic processing feature lets you manage multivendor networks with the extended DHCP and DHCPv6 relay agent.

Related Documentation

- [Extended DHCP Relay Agent Overview](#)
- [Grouping Interfaces with Common DHCP Configurations](#)
- [Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 73](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 75](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 79](#)

Using DHCP Option Information to Selectively Process DHCP Client Traffic

Starting in Junos OS Release 15.1, you can configure the DHCP relay agent to selectively process client traffic. Selective processing uses DHCP or DHCPv6 option information to identify, filter, and process client traffic. To configure DHCPv6 support you use the procedure at the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level.

To configure DHCP relay agent to use option information to selectively process DHCP client traffic:

1. Specify that you want to configure DHCP relay agent support.

```
[edit forwarding-options]  
user@host# edit dhcp-relay
```

2. Specify that you want to use the DHCP option feature to selectively process incoming DHCP traffic.

```
[edit forwarding-options dhcp-relay]  
user@host# edit relay-option
```

3. Specify the DHCP or DHCPv6 option number DHCP relay uses to identify and process the client traffic. You can specify options 60 and 77 for DHCP relay agent, and options 15 and 16 for DHCPv6 relay agent.

```
[edit forwarding-options dhcp-relay relay-option]  
user@host# set option-number option-number
```

For example, to identify traffic that has DHCP option 60 information:

```
[edit forwarding-options dhcp-relay relay-option]  
user@host# set option-number 60
```

4. (Optional) Configure the default action that DHCP relay uses when the incoming client traffic does not satisfy any configured match or partial match criteria.

For example, to configure DHCP relay to drop traffic by default:

```
[edit forwarding-options dhcp-relay relay-option]  
user@host# set default-action drop
```

5. (Optional) Configure an exact match condition that filters the client traffic and specifies the associated action for DHCP relay agent to take.

For example, to select traffic that has an option 60 ASCII string of **video25**, and then forward that traffic to a named local server group:

```
[edit forwarding-options dhcp-relay relay-option]  
user@host# set equals ascii video25 local-server-group servergroup-east-video
```

6. (Optional) Configure a partial match condition that filters the client traffic and specifies the associated action.

For example, to select traffic that has an option 60 hexadecimal string that starts with **766964656F** (left to right), and then forward that traffic without creating a new session:

```
[edit forwarding-options dhcp-relay relay-option]  
user@host# edit starts-with hexadecimal 766964656F forward-only
```

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, you can configure the DHCP relay agent to selectively process client traffic.

Related Documentation

- [DHCP Options and Selective Traffic Processing Overview on page 71](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 75](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 79](#)

Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings

This example shows how to configure DHCP relay agent to use DHCP option strings to selectively identify, filter, and process client traffic.

- [Requirements on page 75](#)
- [Overview on page 75](#)
- [Configuration on page 76](#)
- [Verification on page 78](#)

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or EX Series Switches

Before you configure DHCP relay agent selective processing support, be sure you:

- Configure DHCP relay agent.
See [Extended DHCP Relay Agent Overview](#).
- (Optional) Configure a named DHCP local server group if you want to forward client traffic to a server group.
See [Grouping Interfaces with Common DHCP Configurations](#).

Overview

In this example, you configure DHCP relay agent to use DHCP option strings in client packets to selectively identify, filter, and process client traffic. To configure selective processing, you perform the following procedures:

1. Identify the client traffic—Specify the DHCP option that DHCP relay agent uses to identify the client traffic you want to process. The option you specify matches the option in the client traffic.

2. Configure a default action—Specify the default processing action, which DHCP relay uses for identified client traffic that does not satisfy any configured match criteria.
3. Create match filters and associate an action with each filter—Specify match criteria that filter the client traffic. The criteria can be an exact match or a partial match with the option string in the client traffic. Associate a processing action with each match criterion.

Configuration

To configure DHCP relay agent selective processing based on DHCP option information, perform these tasks:

- [Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings on page 76](#)
- [Results on page 77](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the command into the CLI at the **[edit]** hierarchy level.

```
set forwarding-options dhcp-relay relay-option option-number 60
set forwarding-options dhcp-relay relay-option equals ascii video-gold forward-only
set forwarding-options dhcp-relay relay-option equals ascii video-bronze local-server-group
  servergroup-15
set forwarding-options dhcp-relay relay-option starts-with hexadecimal ffff
  local-server-group servergroup-east
set forwarding-options dhcp-relay relay-option default-action drop
```

Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings

Step-by-Step Procedure

To configure DHCP relay selective processing:

1. Specify that you want to configure DHCP relay agent support.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```
2. Specify the DHCP option that DHCP relay agent uses to identify incoming client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option option-number 60
```
3. Configure a default action, which DHCP relay agent uses when the incoming client traffic does not satisfy any configured match criteria.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option default-action drop
```


4. Configure an exact match condition and associated action that DHCP relay uses to process the identified client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-gold forward-only
```

5. Configure a second exact match condition and associated action that DHCP relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-bronze local-server-group
servergroup-15
```

6. Configure a partial match criteria and associated action that DHCP relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option starts-with hexadecimal ffff local-server-group
servergroup-east
```

Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options]
user@host# show
dhcp-relay {
  relay-option {
    option-number 60;
    equals {
      ascii video-gold {
        forward-only;
      }
    }
    equals {
      ascii video-bronze {
        local-server-group servergroup-15;
      }
    }
    default-action {
      drop;
    }
    starts-with {
      hexadecimal ffff {
        local-server-group servergroup-east;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To verify the status of DHCP relay agent selective traffic processing, perform this task:

- [Verifying the Status of DHCP Relay Agent Selective Traffic Processing on page 78](#)

Verifying the Status of DHCP Relay Agent Selective Traffic Processing

Purpose Verify the DHCP relay agent selective traffic processing status.

Action Display statistics for DHCP relay agent.

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                  30
  Bad hardware address   1
  Bad opcode             1
  Bad options            3
  Invalid server address  5
  No available addresses  1
  No interface match     2
  No routing instance match 9
  No valid local address  4
  Packet too short       2
  Read error             1
  Send error             1
  Option 60              1
  Option 82              2

Messages received:
  BOOTREQUEST            116
  DHCPDECLINE            0
  DHCPDISCOVER           11
  DHCPINFORM             0
  DHCPRELEASE            0
  DHCPREQUEST            105

Messages sent:
  BOOTREPLY              0
  DHCPOFFER              2
  DHCPACK                1
  DHCPNAK                0
  DHCPFORCERENEW         0

Packets forwarded:
  Total                  4
  BOOTREQUEST            2
  BOOTREPLY              2
```

Meaning The **Packets forwarded** field in the **show dhcp relay statistics** command output displays the number of client packets that have been forwarded as a result of the selective traffic processing configuration. In this example, the output indicates the total number of packets

that DHCP relay agent has forwarded, as well as a breakdown for the number of **BOOTREQUEST** and **BOOTREPLY** packets forwarded.

**Related
Documentation**

- [Extended DHCP Relay Agent Overview](#)
- [DHCP Options and Selective Traffic Processing Overview on page 71](#)
- [Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 73](#)
- [Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 79](#)

Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing

This example shows how to configure named interface group-based support for DHCPv6 relay agent selective processing, which uses DHCP option strings to identify, filter, and process client traffic.

This example describes DHCPv6 relay agent configuration—you can configure the related procedure for DHCP relay agent groups at the **[edit forwarding-options dhcp-relay]** hierarchy level. DHCPv6 selective processing supports DHCPv6 options 15 and 16. DHCP selective processing supports option 60 (MX Series routers only) and option 77.

- [Requirements on page 79](#)
- [Overview on page 80](#)
- [Configuration on page 80](#)
- [Verification on page 82](#)

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or PTX Series Packet Transport Routers

Before you configure DHCPv6 relay agent selective processing support, be sure you:

- Configure DHCPv6 relay agent.

See [Extended DHCP Relay Agent Overview](#) and [DHCPv6 Relay Agent Overview](#).

- Configure the DHCPv6 named interface groups used for the configuration.

See [Grouping Interfaces with Common DHCP Configurations](#).

- Configure the DHCPv6 server groups used for the processing actions.

See [Grouping Interfaces with Common DHCP Configurations](#).

Overview

In this example, you configure group-level DHCPv6 relay agent named interface support for selective processing of client packets based on DHCPv6 option strings. To configure selective processing, you perform the following procedures:

1. Identify the client traffic—Specify the DHCPv6 option that DHCPv6 relay agent uses to identify the client traffic you want to process. The DHCPv6 option you specify matches the option in the client traffic.
2. Configure the default action—Specify the default processing action, which DHCPv6 relay uses for identified client traffic that does not satisfy any configured match criteria.
3. Create match filters and associate an action with each filter—Specify match criteria that filters the client traffic. The criteria can be an exact match or a partial match with the DHCPv6 option string in the client traffic. Associate a processing action with each match criteria.

Configuration

To configure group-level DHCPv6 relay agent selective processing based on DHCPv6 option information, perform these tasks:

- [Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings on page 80](#)
- [Results on page 81](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the command into the CLI at the **[edit]** hierarchy level. The quick configuration assumes that the named interface group and the DHCP server groups have been previously configured.

```
set forwarding-options dhcp-relay dhcpv6 group groupv6-east-27
set forwarding-options dhcp-relay dhcpv6 relay-option option-number 15
set forwarding-options dhcp-relay dhcpv6 relay-option equals ascii triple-gold
  relay-server-group relayserver-triple-8
set forwarding-options dhcp-relay dhcpv6 relay-option equals ascii triple-silver
  relay-server-group relayserver-triple-23
set forwarding-options dhcp-relay dhcpv6 relay-option starts-with ascii single
  relay-server-group relayserver-1-aa
set forwarding-options dhcp-relay dhcpv6 relay-option default-action drop
```

[Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings](#)

Step-by-Step Procedure

This procedure assumes that you have previously created the named interface group and the DHCPv6 server groups. To configure DHCPv6 relay group-level selective processing:

1. Specify that you want to configure DHCPv6 relay agent support.

```
[edit forwarding-options]
user@host# edit dhcp-relay dhcpv6
```

2. Specify that you want to configure group-level DHCPv6 relay agent support.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group groupv6-east-27
```

3. Specify the DHCPv6 option number that DHCPv6 relay agent uses to identify incoming client traffic.

```
[edit forwarding-options dhcp-relay dhcpv6 group groupv6-east-27]
user@host# set relay-option option-number 15
```

4. Configure the default action, which DHCPv6 relay agent uses when the incoming client traffic does not satisfy any configured match criteria.

```
[edit forwarding-options dhcp-relay dhcpv6 group groupv6-east-27]
user@host# set relay-option default-action relay-server-group relayserver-def-4
```

5. Configure an exact match condition and associated action that DHCPv6 relay uses to process the identified client traffic.

```
[edit forwarding-options dhcp-relay dhcpv6 group groupv6-east-27]
user@host# set relay-option equals ascii triple-gold relay-server-group
relayserver-triple-8
```

6. Configure a second exact match condition and associated action that DHCPv6 relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay dhcpv6 group groupv6-east-27]
user@host# set relay-option equals ascii triple-silver relay-server-group
relayserver-triple-23
```

7. Configure a partial match criteria and associated action that DHCPv6 relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay dhcpv6 group groupv6-east-27]
user@host# set relay-option starts-with ascii single relay-server-group
relayserver-1-aa
```

Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options dhcp]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
dhcpv6 {
  group test-1 {
    relay-option {
      option-number 15;
    }
  }
}
```

```
equals {
  ascii triple-gold {
    relay-server-group relayserver-triple-8;
  }
  ascii triple-silver {
    relay-server-group relayserver-triple-23;
  }
}
default-action {
  relay-server-group relayserver-def-4;
}
starts-with {
  ascii single {
    relay-server-group relayserver-1-aa;
  }
}
}
interface ge-1/0/0.0 upto ge-1/1/0.0;
}
server-group {
  relayserver-1-aa;
  relayserver-triple-8;
  relayserver-triple-23;
  relayserver-def-4;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To verify the status of DHCPv6 relay agent selective traffic processing, perform this task:

- [Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing on page 82](#)

Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing

Purpose Verify the DHCPv6 relay agent selective traffic processing status.

Action Display statistics for DHCPv6 relay agent.

```
user@host> show dhcpv6 relay statistics
```

```
DHCPv6 Packets dropped:
```

```
  Total                               0
```

```
Messages received:
```

```
  DHCPV6_DECLINE                      0
  DHCPV6_SOLICIT                      10
  DHCPV6_INFORMATION_REQUEST          0
  DHCPV6_RELEASE                      0
  DHCPV6_REQUEST                      10
  DHCPV6_CONFIRM                      0
  DHCPV6_RENEW                        0
  DHCPV6_REBIND                       0
  DHCPV6_RELAY_REPL                   0
```

```
Messages sent:
```

```
  DHCPV6_ADVERTISE                    0
  DHCPV6_REPLY                        0
  DHCPV6_RECONFIGURE                  0
  DHCPV6_RELAY_FORW                   0
```

```
Packets forwarded:
```

```
  Total                               4
  FWD REQUEST                         2
  FWD REPLY                           2
```

Meaning The **Packets forwarded** field in the **show dhcpv6 relay statistics** command output displays the number of client packets that have been forwarded as a result of the selective traffic processing configuration. In this example, the output indicates the total number of packets that DHCPv6 relay agent has forwarded, as well as a breakdown for the number of **FWD REQUEST** and **FWD REPLY** packets forwarded.

Related Documentation

- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*
- [DHCP Options and Selective Traffic Processing Overview on page 71](#)
- [Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 73](#)
- *Grouping Interfaces with Common DHCP Configurations*
- *Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings*
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 75](#)

CHAPTER 9

Configuring High Availability in the DHCP Access Network

- [DHCP Liveness Detection Overview on page 85](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87](#)
- [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)
- [Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92](#)
- [Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94](#)
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)
- [High Availability Using Unified ISSU in the DHCP Access Network on page 106](#)
- [Graceful Routing Engine Switchover for DHCP on page 106](#)
- [Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 107](#)
- [Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 108](#)

DHCP Liveness Detection Overview

Unlike PPP, DHCP does not define a native keepalive mechanism as part of either the DHCPv4 or DHCPv6 protocols. Without a keepalive mechanism, DHCP local server, DHCP relay, and DHCP relay proxy are unable to quickly detect if any of them has lost connectivity with a subscriber or a DHCP client. Instead, they must rely on standard DHCP subscriber session or DHCP client session termination messages.

DHCP clients often do not send DHCP release messages before exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection

protocol, if a given subscriber (or client) fails to respond to a configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

DHCP liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

Using DHCP liveness detection, IP sessions are acted upon as soon as liveness detection checks fail. This faster response time serves to:

- Provide more accurate time-based accounting of subscriber (or DHCP client) sessions.
- Better preserve router (switch) resources.
- Help to reduce the window of vulnerability to some security attacks.

Examples of liveness detection protocols include Bidirectional Forwarding Detection (BFD) for both DHCPv4 and DHCPv6 subscribers, IPv4 Address Resolution Protocol (ARP) for DHCPv4 subscribers, and IPv6 Neighbor Unreachability Detection (NUD) using Neighbor Discovery (ND) packets for DHCPv6 subscribers.

Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection. In earlier releases, only BFD is supported for all platforms.

The two liveness detection methods are mutually exclusive.

When configuring BFD liveness detection, keep the following in mind:

- You can configure liveness detection for both DHCP local server and DHCP relay.
- You can configure DHCPv4 and DHCPv6 liveness detection either globally or per DHCPv4 or DHCPv6 group.
- DHCPv4 or DHCPv6 subscriber access clients that do not support BFD are not affected by the liveness detection configuration. These clients can continue to access the network (after they are validated) even if BFD liveness detection is enabled on the router (or switch).
- When configured, DHCPv4 or DHCPv6 initiates liveness detection checks for clients that support BFD when those clients enter a bound state.
- After protocol-specific messages are initiated for a BFD client, they are periodically sent to the subscriber (or client) IP address of the client and responses to those liveness detection requests are expected within a configured amount of time.
- If liveness detection responses are not received from clients that support BFD within the configured amount of time for a configured number of consecutive attempts, the

liveness detection check is deemed to have failed. A configured failure action to clear the client binding is applied.

- The only failure action supported for Layer 2 Liveness detection is **clear-binding**.

When configuring DHCP ARP and ND Layer 2 liveness detection on MX Series, keep the following in mind:

- You can configure liveness detection for both DHCP local server and DHCP relay.
- You can configure DHCPv4 and DHCPv6 ARP and ND liveness detection globally, per DHCPv4 or DHCPv6 group, and per dual-stack group.
- ARP/ND liveness detection applies only to DHCP clients that:
 - Are directly connected over dynamic VLANs.
 - Have permanent Layer 2 entries.
- DHCPv6 clients must have a unique source MAC address and link-local address. Only single liveness detection entry is used for all IPv6 addresses associated with a specific client session.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection.

Related Documentation

- [Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87](#)
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)

Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD

You can configure liveness detection with Bidirectional Forwarding Detection (BFD) for DHCP subscriber IP sessions or DHCP client IP sessions to check the connectivity of DHCP relay clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit liveness-detection
```



NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the **liveness-detection** statement, and any subsequent configuration statements, at the `[edit forwarding-options dhcp-relay dhcpv6]` or `[edit forwarding-options dhcp-relay dhcpv6 group group-name]` hierarchy level.

2. (Optional) Specify that you want to use DHCP relay proxy mode.

```
[edit forwarding-options dhcp-relay group group-name]
user@host# set overrides proxy-mode
```

3. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit method
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]
user@host# edit method
```

4. Specify the liveness detection method that you want DHCP to use.



NOTE: In releases earlier than Junos OS Release 17.4R1, the only method supported for liveness detection on all platforms is BFD.

Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection. The two liveness detection methods are mutually exclusive. See [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets](#) for information about configuring ARP and ND Layer 2 liveness detection.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method]
user@host# edit **bfd**

5. Configure the liveness detection method as desired.

See [“Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients” on page 89](#) for an example of how to globally configure DHCP relay liveness detection with BFD.

6. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit **failure-action** *action*
- For DHCP group configuration:
[edit forwarding-options dhcp-relay group *group-name* liveness-detection]
user@host# edit **failure-action** *action*

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection.

Related Documentation

- [Extended DHCP Relay Agent Overview](#)
- [DHCP Liveness Detection Overview on page 85](#)
- [Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92](#)
- [Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94](#)
- [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

This example shows how to configure liveness detection for DHCP relay agent subscribers using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

- [Requirements on page 89](#)
- [Overview on page 90](#)
- [Configuration on page 90](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks MX Series routers.
- Junos OS Release 12.1 or later

Before you begin:

- Configure DHCP relay agent. See *Extended DHCP Relay Agent Overview*.

Overview

In this example, you configure liveness detection for DHCP relay agent subscribers by completing the following operations:

1. Enable liveness detection globally for DHCP relay subscribers.
2. Specify BFD as the liveness detection method for all dynamically created DHCP relay subscribers.
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router takes when a liveness detection failure occurs.



NOTE: This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the `liveness-detection` statement, and any subsequent configuration statements, at the `[edit forwarding-options dhcp-relay dhcpv6]` or `[edit forwarding-options dhcp-relay dhcpv6 group group-name]` hierarchy level.

Configuration

Step-by-Step Procedure

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.
`[edit forwarding-options dhcp-relay]`
user@host# **edit** `liveness-detection`
2. Specify that you want to configure the liveness detection method.
`[edit forwarding-options dhcp-relay liveness-detection]`
user@host# **edit** `method`
3. Specify BFD as the liveness detection method that you want DHCP to use.
`[edit forwarding-options dhcp-relay liveness-detection method]`
user@host# **edit** `bfd`

4. Configure the detection time threshold (in milliseconds) at which a trap is produced.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **detection-time threshold** 50000
5. Configure the time (in milliseconds) for which BFD holds a session up notification.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **holddown-interval** 50
6. Configure the BFD minimum transmit and receive interval (in milliseconds).
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **minimum-interval** 45000
7. Configure the minimum receive interval (in milliseconds).
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **minimum-receive-interval** 60000
8. Configure a multiplier value for the detection time.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **multiplier** 100
9. Disable the ability for BFD interval timers to change or adapt to network situations.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **no-adaptation**
10. Configure the BFD session mode.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **session-mode** automatic
11. Configure the threshold and minimum interval for the BFD transmit interval.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **transmit-interval threshold** 60000 **minimum-interval** 45000
12. Configure the BFD protocol version you want to detect.
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set **version** automatic
13. Configure the action the router takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit **failure-action** action

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group frankfurt.

```
[edit]
user@host# show forwarding-options
dhcp-relay {
  liveness-detection {
    failure-action clear-binding-if-interface-up;
    method {
      bfd {
        version automatic;
        minimum-interval 45000;
        minimum-receive-interval 60000;
        multiplier 100;
        no-adaptation;
        transmit-interval {
          minimum-interval 45000;
          threshold 60000;
        }
        detection-time {
          threshold 50000;
        }
        session-mode automatic;
        holddown-interval 50;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Extended DHCP Relay Agent Overview](#)
 - [DHCP Liveness Detection Overview on page 85](#)
 - [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87](#)

Configuring Detection of DHCP Local Server Client Connectivity with BFD

You can configure liveness detection with Bidirectional Forwarding Detection (BFD) for DHCP subscriber IP sessions or DHCP client IP sessions to check the connectivity of DHCP local server clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



NOTE: You can also configure DHCP liveness detection for DHCP relay.

To configure liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit system services dhcp-local-server]
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name]
user@host# edit liveness-detection
```



NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the **liveness-detection** statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

2. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit method
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit method
```

3. Specify the liveness detection method that you want DHCP to use.



NOTE: In releases earlier than Junos OS Release 17.4R1, the only method supported for liveness detection on all platforms is BFD.

Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection. The two liveness detection methods are mutually exclusive. See [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets](#) for information about configuring ARP and ND Layer 2 liveness detection.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection method]
user@host# edit bfd
```

4. Configure the liveness detection method as desired.

See [“Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients” on page 94](#) for an example of how to configure DHCPv4 groups for DHCP local server liveness detection with BFD.

5. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit failure-action action
```

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the use of ARP packets for DHCPv4 and ND packets for DHCPv6 is supported on MX Series routers for Layer 2 liveness detection in addition to BFD liveness detection.

Related Documentation

- [DHCP Liveness Detection Overview on page 85](#)
- [Extended DHCP Local Server Overview](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87](#)
- [Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94](#)
- [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

This example shows how to configure group liveness detection for DHCP local server subscribers or DHCP clients using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

- [Requirements on page 94](#)
- [Overview on page 95](#)
- [Configuration on page 95](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks MX Series routers

- Juniper Networks EX Series switches
- Junos OS Release 12.1 or later

Before you begin:

- Configure DHCP local server. See *Extended DHCP Local Server Overview*.

Overview

In this example, you configure group liveness detection for DHCP local server subscribers (clients) by completing the following operations:

1. Enable liveness detection for DHCP local server subscriber (or DHCP client) groups.
2. Specify BFD as the liveness detection method for all dynamically created DHCP local server subscribers (clients).
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router (switch) takes when a liveness detection failure occurs.



NOTE: This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the `liveness-detection` statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

Configuration

Step-by-Step Procedure

To configure group liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

```
[edit system services dhcp-local-server ]
user@host# edit liveness-detection
```
2. Specify that you want to configure liveness detection for a specific DHCP local server group.

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit group local_group_1
```
3. Specify that you want to configure the liveness detection method.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit method
```

4. Specify BFD as the liveness detection method that you want DHCP to use.
[edit system services dhcp-local-server group local_group_1 liveness-detection method]
user@host# **edit bfd**
5. Configure the detection time threshold (in milliseconds) at which a trap is produced.
[edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
user@host# **set detection-time threshold 30000**
6. Configure the time (in milliseconds) for which BFD holds a session up notification.
[edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
user@host# **set holddown-interval 50**
7. Configure the BFD minimum transmit and receive interval (in milliseconds).



NOTE: You do not need to configure the BFD minimum transmit and receive interval if you configure the minimum-interval for the BFD transmit-interval statement and the minimum-receive-interval.

[edit system services dhcp-local-servergroup local_group_1 liveness-detection method bfd]
user@host# **set minimum-interval 45000**

8. Configure the minimum receive interval (in milliseconds).



NOTE: You do not need to configure the BFD minimum receive interval if you configure the BFD minimum transmit and receive interval.

[edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
user@host# **set minimum-receive-interval 60000**

9. Configure a multiplier value for the detection time.
[edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
user@host# **set multiplier 100**
10. Disable the ability for BFD interval timers to change or adapt to network situations.
[edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
user@host# **set no-adaptation**

11. Configure the BFD session mode.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set session-mode automatic
```

12. Configure the threshold and minimum interval for the BFD transmit interval.



NOTE: You do not need to configure the transmit interval values if you have already configured the minimum transmit and receive interval for BFD.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set transmit-interval threshold 60000 minimum-interval 45000
```

13. Configure the BFD protocol version you want to detect.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set version automatic
```

14. Configure the action the router (switch) takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit failure-action action
```

Results From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show system
services {
  dhcp-local-server {
    group local_group_1 {
      liveness-detection {
        failure-action clear-binding-if-interface-up;
        method {
          bfd {
            version automatic;
            minimum-interval 45000;
            minimum-receive-interval 60000;
            multiplier 100;
            no-adaptation;
            transmit-interval {
              minimum-interval 45000;
              threshold 60000;
            }
          }
        }
      }
    }
  }
}
```

```
    }  
    detection-time {  
        threshold 30000;  
    }  
    session-mode automatic;  
    holddown-interval 50;  
  }  
}  
}  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related
Documentation**

- [Extended DHCP Local Server Overview](#)
- [DHCP Liveness Detection Overview on page 85](#)
- [Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92](#)

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

- [How DHCP Liveness Detection with ARP and Neighbor Discovery Packets Works on page 98](#)
- [Configuring BNG Detection of DHCP Local Server Client Connectivity with ARP and ND Packets on page 101](#)
- [Configuring BNG Detection of DHCP Relay Client Connectivity with ARP and ND Packets on page 103](#)
- [Configuring DHCP Host Detection of Client Connectivity with ARP and ND Packets on page 105](#)

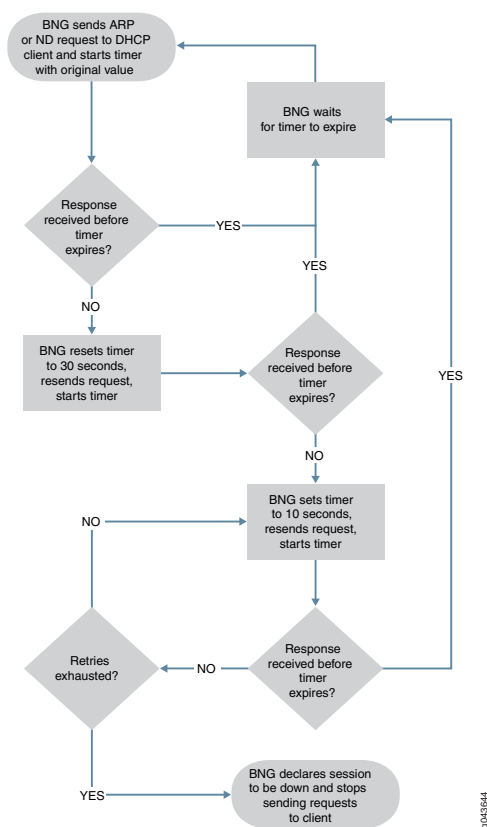
How DHCP Liveness Detection with ARP and Neighbor Discovery Packets Works

Starting in Junos OS Release 17.4R1, you can configure liveness detection using IPv4 Address Resolution Protocol (ARP) for DHCPv4 clients and IPv6 Neighbor Unreachability Detection for DHCPv6 clients. This Layer 2 liveness detection offers separate mechanisms for the DHCP client host and for the router acting as a broadband network gateway (BNG) to determine the validity and state of the DHCP client sessions. These mechanisms are referred to as the *send* functionality and the *receive* functionality. You can configure Layer 2 liveness detection for DHCP local server and DHCP relay clients.

Send Functionality

The BNG uses the send functionality to conduct a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients to determine the validity and state of the DHCP client session, and to clean up inactive sessions. [Figure 6 on page 99](#) illustrates the send functionality.

Figure 6: Layer 2 Liveness Detection Send Behavior Flow



1. The BNG sends request packets to the each DHCP client at a configurable interval, then waits for a response. The BNG retries the requests when it does not receive a timely response. It sends ARP requests for DHCPv4 clients and Neighbor Discovery (ND) requests for DHCPv6 clients.
2. If the BNG receives a response from the client before the interval times out, it waits for the timer to expire and then sends another request to that client.
3. If the BNG does not receive a response before the interval times out, it sets the timer to 30 seconds and sends another request. This is the first retry attempt; the timer is not configurable.
4. If the BNG receives a response from the client before the timer expires, then the BNG waits for the timer to run down, resets it to the original, configurable value, sends another request, and starts the timer.
5. If the 30-second timer expires before a response is received, the BNG sets the timer to 10 seconds and sends another request. This timer value is not configurable.

6. If the BNG receives a response from the client before the timer expires, then the BNG waits for the timer to run down, resets it to the original, configurable value, sends another request, and starts the timer.
7. If the BNG does not receive a response within the 10-second interval, it sends another request and starts the 10-second timer again. The BNG continues to send requests at 10-second intervals until it receives a response from the client before the interval times out or it exhausts the number of retry attempts.

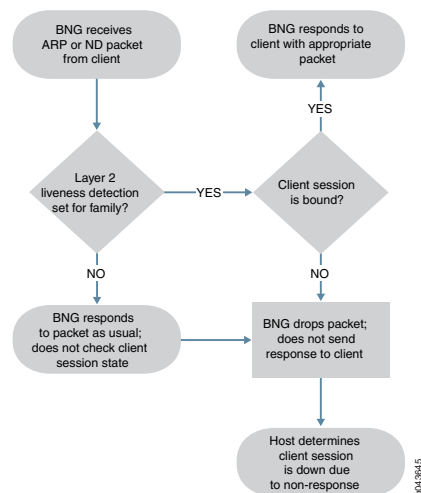
The first retry attempt uses the 30-second interval. Subsequent retries occur at 10-second intervals. The number of possible 10-second retries is therefore the total number minus 1. For example, if you configure 5 retries, there is one 30-second retry and up to four 10-second retries.

8. If the BNG never sends a response from a client within the interval before the retries are exhausted, then the liveness detection check fails and the clear-binding failure action is implemented. The client session is cleared.

Receive Functionality

The receive functionality enables a DHCP client host to determine the state of the DHCPv4 or DHCPv6 client session from the perspective of a BNG. The BNG conducts a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients when it receives ARP or ND packets. [Figure 7 on page 100](#) illustrates the receive functionality.

Figure 7: Layer 2 Liveness Detection Receive Behavior Flow



When the BNG receives either of these packets, it does the following:

1. Checks whether Layer 2 liveness detection for subscriber management is enabled globally for the relevant address family, inet or inet6.
2. If Layer 2 liveness detection is not enabled, then the BNG responds as usual to the received packets without checking the state of the client session.

3. If liveness detection is enabled for the family, then the BNG checks whether the client session is still in the bound state.
4. If the client session is bound, the BNG responds to the client with the appropriate ARP or ND packet.
5. If the session is not bound, the BNG drops the received packet. It does not send an ARP or ND response packet to the host, enabling the host to determine that the BNG considers the session to be down.

The usefulness of the receive functionality depends on the ability of the DHCP client host to reclaim resources from the stale client based on the absence of a response packet from the BNG for an unbound client session. If this capability requires a change in the client implementation, you may want to use the send functionality.

Configuring BNG Detection of DHCP Local Server Client Connectivity with ARP and ND Packets

This procedure shows you how to configure the send functionality of Layer 2 liveness detection using IPv4 Address Resolution Protocol (ARP) for DHCPv4 clients and IPv6 Neighbor Unreachability Detection for DHCPv6 clients to check the connectivity of DHCP local server clients.

The send functionality enables the BNG to determine whether a client session is down based on a lack of response from the DHCP client to the ARP or ND request packets it sends to the client.



NOTE: DHCP liveness detection can also be configured using Bidirectional Forwarding Detection (BFD). BFD liveness detection and ARP/ND liveness detection are mutually exclusive.

To configure the send functionality for DHCPv4 local server liveness detection:

1. Specify that you want to configure the liveness detection method.

- For DHCPv4 global configuration:

```
[edit system services dhcp-local-server]
user@host# edit liveness-detection method
```

- For DHCPv4 group configuration:

```
[edit system services dhcp-local-server]
user@host# edit group group-name liveness-detection method
```

- For DHCPv4 dual-stack-group configuration:

```
[edit system services dhcp-local-server]
user@host# edit dual-stack-group dual-stack-group-name liveness-detection
method
```

2. Specify the Layer 2 liveness detection method.

- For DHCPv4 global configuration:
[edit system services dhcp-local-server liveness-detection *method*]
user@host# set layer2-liveness-detection
 - For DHCPv4 group configuration:
[edit system services dhcp-local-server group *group-name* liveness-detection *method*]
user@host# set layer2-liveness-detection
 - For DHCPv4 dual-stack-group configuration:
[edit system services dhcp-local-server dual-stack-group *dual-stack-group-name* liveness-detection *method*]
user@host# set layer2-liveness-detection
3. (Optional) Configure the number of retry attempts and the interval timer.
- For DHCPv4 global configuration:
[edit system services dhcp-local-server liveness-detection *method*]
user@host# edit layer2-liveness-detection
user@host# set max-consecutive-retries *number*
user@host# set transmit-interval *seconds*
 - For DHCPv4 group configuration:
[edit system services dhcp-local-server group *group-name* liveness-detection *method*]
user@host# edit layer2-liveness-detection
user@host# set max-consecutive-retries *number*
user@host# set transmit-interval *seconds*
 - For DHCPv4 dual-stack-group configuration:
[edit system services dhcp-local-server dual-stack-group *dual-stack-group-name* liveness-detection *method*]
user@host# edit layer2-liveness-detection
user@host# set max-consecutive-retries *number*
user@host# set transmit-interval *seconds*

To configure the send functionality for DHCPv6 local server liveness detection:

1. Specify that you want to configure the liveness detection method.
 - For DHCPv6 global configuration:
[edit system services dhcp-local-server dhcpv6]
user@host# edit liveness-detection *method*
 - For DHCPv6 group configuration:
[edit system services dhcp-local-server dhcpv6]
user@host# edit group *group-name* liveness-detection *method*
2. Specify the Layer 2 liveness detection method.
 - For DHCPv6 global configuration:
[edit system services dhcp-local-server dhcpv6 liveness-detection *method*]
user@host# set layer2-liveness-detection
 - For DHCPv6 group configuration:

```
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection
method]
user@host# set layer2-liveness-detection
```

3. (Optional) Configure the number of retry attempts and the interval timer.

- For DHCPv6 global configuration:

```
[edit system services dhcp-local-server dhcpv6 liveness-detection method]
user@host# edit layer2-liveness-detection
user@host# set max-consecutive-retries number
user@host# set transmit-interval seconds
```

- For DHCPv6 group configuration:

```
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection
method]
user@host# edit layer2-liveness-detection
user@host# set max-consecutive-retries number
user@host# set transmit-interval seconds
```

Configuring BNG Detection of DHCP Relay Client Connectivity with ARP and ND Packets

This procedure shows you how to configure the send functionality of Layer 2 liveness detection using IPv4 Address Resolution Protocol (ARP) for DHCPv4 clients and IPv6 Neighbor Unreachability Detection for DHCPv6 clients to check the connectivity of DHCP relay clients.

The send functionality enables the BNG to determine whether a client session is down based on a lack of response from the DHCP client to the ARP or ND request packets it sends to the client.



NOTE: DHCP liveness detection can also be configured using Bidirectional Forwarding Detection (BFD). BFD liveness detection and ARP/ND liveness detection are mutually exclusive.

To configure the send functionality for DHCPv4 relay liveness detection:

1. Specify that you want to configure the liveness detection method.

- For DHCPv4 global configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit liveness-detection method
```

- For DHCPv4 group configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit group group-name liveness-detection method
```

- For DHCPv4 dual-stack-group configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit dual-stack-group dual-stack-group-name liveness-detection
method
```

2. Specify the Layer 2 liveness detection method.

- For DHCPv4 global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection method]  
user@host# set layer2-liveness-detection
```
 - For DHCPv4 group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method]  
user@host# set layer2-liveness-detection
```
 - For DHCPv4 dual-stack-group configuration:

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name  
liveness-detection method]  
user@host# set layer2-liveness-detection
```
3. (Optional) Configure the number of retry attempts and the interval timer.
- For DHCPv4 global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection method]  
user@host# edit layer2-liveness-detection  
user@host# set max-consecutive-retries number  
user@host# set transmit-interval seconds
```
 - For DHCPv4 group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method]  
user@host# edit layer2-liveness-detection  
user@host# set max-consecutive-retries number  
user@host# set transmit-interval seconds
```
 - For DHCPv4 dual-stack-group configuration:

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name  
liveness-detection method]  
user@host# edit layer2-liveness-detection  
user@host# set max-consecutive-retries number  
user@host# set transmit-interval seconds
```

To configure the send functionality for DHCPv6 relay liveness detection:

1. Specify that you want to configure the liveness detection method.
 - For DHCPv6 global configuration:

```
[edit forwarding-options dhcp-relay dhcpv6]  
user@host# edit liveness-detection method
```
 - For DHCPv6 group configuration:

```
[edit forwarding-options dhcp-relay dhcpv6]  
user@host# edit group group-name liveness-detection method
```
2. Specify the Layer 2 liveness detection method.
 - For DHCPv6 global configuration:

```
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method]  
user@host# set layer2-liveness-detection
```
 - For DHCPv6 group configuration:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection
method]
```

```
user@host# set layer2-liveness-detection
```

3. (Optional) Configure the number of retry attempts and the interval timer.

- For DHCPv6 global configuration:

```
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method]
```

```
user@host# edit layer2-liveness-detection
```

```
user@host# set max-consecutive-retries number
```

```
user@host# set transmit-interval seconds
```

- For DHCPv6 group configuration:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection
method]
```

```
user@host# edit layer2-liveness-detection
```

```
user@host# set max-consecutive-retries number
```

```
user@host# set transmit-interval seconds
```

Configuring DHCP Host Detection of Client Connectivity with ARP and ND Packets

This procedure shows you how to configure the receive functionality of Layer 2 liveness detection using IPv4 Address Resolution Protocol (ARP) for DHCPv4 clients and IPv6 Neighbor Unreachability Detection for DHCPv6 clients to check the connectivity of DHCP local server clients.

The receive functionality enables the DHCP client host to determine whether a client session is down based on a lack of response from the BNG to the ARP or ND packets it sends to the BNG. You configure the receive functionality globally for DHCP per address family as an override to the global subscriber management configuration.

1. Enable Layer 2 liveness detection globally per address family.

- For DHCPv4:

```
[edit system services subscriber-management overrides]
```

```
user@host# set interfaces family inet layer2-liveness-detection
```

- For DHCPv6:

```
[edit system services subscriber-management overrides]
```

```
user@host# set interfaces family inet6 layer2-liveness-detection
```

See Also • [DHCP Liveness Detection Overview on page 85](#)

High Availability Using Unified ISSU in the DHCP Access Network

Starting in Junos OS Release 14.1, the unified in-service software upgrade (unified ISSU) feature supports the DHCP access model used by subscriber management. This support ensures that the router preserves active DHCP subscriber sessions and session services after a unified ISSU has completed.

See *Getting Started with Unified In-Service Software Upgrade* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

Unified ISSU supports the subscriber management DHCP access model, which includes DHCP local server, DHCPv6 local server, DHCP relay, and DHCP relay proxy.

Accounting, filter, and class-of-service (CoS) statistics for DHCP subscribers are preserved after a unified ISSU on MPC/MIC interfaces on MX Series routers.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, the unified in-service software upgrade (unified ISSU) feature supports the DHCP access model used by subscriber management.

- Related Documentation**
- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 111](#)
 - [Unified ISSU System Requirements](#)

Graceful Routing Engine Switchover for DHCP

For EX Series switches, only extended DHCP local server maintains the state of active DHCP client leases. The DHCP local server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can configure dynamic profile and authentication support on a global basis or for a specific group of interfaces. The extended DHCP local server also supports the use of Junos address-assignment pools or external authorities, such as RADIUS, to provide the client address and configuration information.

For MX Series routers, the extended DHCP local server and the DHCP relay agent applications both maintain the state of active DHCP client leases in the session database. The extended DHCP application can recover this state if the DHCP process fails or is manually restarted, thus preventing the loss of active DHCP clients in either of these circumstances. However, the state of active DHCP client leases is lost if a power failure occurs or if the kernel stops operating (for example, when the router is reloaded) on a single Routing Engine.

You can enable graceful switchover support on both EX Series switches and MX Series routers. To enable graceful switchover support for the extended DHCP local server or extended DHCP relay agent on a switch, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To enable graceful Routing Engine switchover support on MX Series routers, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. You cannot disable graceful Routing Engine switchover support for the extended DHCP application when the router is configured to support graceful Routing Engine switchover.

For more information about using graceful Routing Engine switchover, see *Understanding Graceful Routing Engine Switchover*.

- Related Documentation**
- [Extended DHCP Local Server Overview](#)
 - [Extended DHCP Relay Agent Overview](#)
 - [High Availability Using Unified ISSU in the PPP Access Network on page 139](#)

Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to wait 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover (GRES) takes place.

- [Benefits of Delaying Removal of Access Routes and Access-Internal Routes on page 107](#)
- [Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes on page 108](#)
- [Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes on page 108](#)

Benefits of Delaying Removal of Access Routes and Access-Internal Routes

The 3-minute delay in removing access routes and access-internal routes after a graceful Routing Engine switchover provides sufficient time for the DHCP client process (jdhcpd), PPP client process (jpppd), or routing protocol process (rpd) to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

Configuring the router to delay removal of access routes and access-internal routes after a graceful Routing Engine switchover has the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine
- Prevents loss of subscriber traffic due to unavailable routes

Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale access routes and access-internal routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the `jdhcpd` or `jpppd` client process to reinstall all of the subscriber routes.

Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process (`rpd`) immediately purges the stale access routes and access-internal routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

- Related Documentation**
- [Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 108](#)
 - [Access and Access-Internal Routes for Subscriber Management on page 31](#)
 - [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover

In subscriber networks configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to delay for 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover takes place.

To configure the router to delay removal (flushing) of access-routes and access-internal routes after a graceful Routing Engine switchover:

1. Specify that you want to configure subscriber management.

```
[edit system services]  
user@host# edit subscriber-management
```

2. Configure the router to wait 180 seconds before removing access-routes and access-internal routes after a graceful Routing Engine switchover.

```
[edit system services subscriber-management]  
user@host# set gres-route-flush-delay
```


- Related Documentation**
- [Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 107](#)
 - [Access and Access-Internal Routes for Subscriber Management on page 31](#)
 - [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

CHAPTER 10

Monitoring and Managing DHCP for Subscriber Access

- [Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 111](#)
- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 111](#)

Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers

Purpose View configuration information for access routes and access-internal routes on DHCP subscribers. The access-internal routes are those that are automatically installed when a client profile is instantiated.

- Action**
- To display extensive information about access routes and access-internal routes:
`user@host>show route extensive`
 - To display the configuration for access routes:
`user@host>show route protocol access`
 - To display the configuration for access-internal routes:
`user@host> show route protocol access-internal`

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane quiescing has not started). The second example shows that unified ISSU is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

`user@host> show system subscriber-management summary`

General:

Graceful Restart	Enabled
Mastership	Master
Database	Available
Chassisd ISSU State	IDLE
ISSU State	IDLE
ISSU Wait	0

user@host> [show system subscriber-management summary](#)

General:

Graceful Restart	Enabled
Mastership	Master
Database	Available
Chassisd ISSU State	DAEMON_ISSU_PREPARE
ISSU State	PREPARE
ISSU Wait	198

**Related
Documentation**

- [High Availability Using Unified ISSU in the PPP Access Network on page 139](#)
- [High Availability Using Unified ISSU in the DHCP Access Network on page 106](#)
- [High Availability Using Unified ISSU in the L2TP Access Network on page 297](#)
- *Getting Started with Unified In-Service Software Upgrade*

PART 3

Configuring the PPP Access Network

- [Configuring PPP for Subscriber Access on page 115](#)
- [Applying RADIUS Route Attributes to Subscribers or Access Networks on page 123](#)
- [Configuring Authentication for PPP on page 125](#)
- [Configuring PPP Network Control Protocol Negotiation on page 129](#)
- [Configuring High Availability in the PPP Access Network on page 139](#)
- [Monitoring and Managing PPP for Subscriber Access on page 143](#)

Configuring PPP for Subscriber Access

- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- [Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests on page 116](#)
- [Configuring Dynamic Profiles for PPP on page 118](#)
- [Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges on page 120](#)
- [Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121](#)
- [Example: Minimum PPPoE Dynamic Profile on page 121](#)

Dynamic Profiles for PPP Subscriber Interfaces Overview

Subscriber management PPP support enables you to create and attach dynamic profiles for PPP subscriber interfaces. When the PPP subscriber logs in, the router instantiates the specified dynamic profile and then applies the attributes defined in the profile to the interface.

Dynamic profiles are used for both static and dynamic PPP interfaces. For static PPP interfaces, you use the CLI to attach dynamic profiles, which specify PPP options. For dynamic PPP interfaces, the dynamic profile creates the interface, including the PPP options.



NOTE: Dynamically created interfaces are supported only on PPPoE interfaces.

Unlike traditional PPP support, subscriber management does not allow bi-directional PPP authentication—authentication is performed only by the router, never by the remote peer. The router's AAA process manages authentication and address assignment for subscriber management. When you configure PPP options for a dynamic profile, you can configure either Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP) authentication, and you can control the order in which the router negotiates the CHAP and PAP protocols. In addition, for CHAP authentication, you can modify the default length of the CHAP challenge message. Other PPP options, which are either commonly used or mandatory for a traditional PPP interface configuration, are not supported in subscriber management dynamic profiles.

Related Documentation

- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121](#)
- [Verifying and Managing PPP Configuration for Subscriber Management on page 143](#)
- [Example: Minimum PPPoE Dynamic Profile on page 121](#)

Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests

On MX Series routers with Modular Port Concentrators/Modular Interface Cards (MPCs/MICs), the Packet Forwarding Engine on an MPC/MIC processes and responds to Link Control Protocol (LCP) Echo-Request packets that the PPP subscriber (client) initiates and sends to the router. LCP Echo-Request packets and LCP Echo-Reply packets are part of the PPP keepalive mechanism that helps determine whether a link is functioning properly.

Previously, LCP Echo-Request packets and LCP Echo-Reply packets were handled on an MX Series router by the Routing Engine. The mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine is referred to as *PPP fast keepalives*.

- [Benefits of PPP Fast Keepalives on page 116](#)
- [How PPP Fast Keepalive Processing Works on page 116](#)
- [Statistics Display for PPP Fast Keepalive on page 117](#)
- [Effect of Changing the Forwarding Class Configuration on page 117](#)
- [Preventing the Validation of PPP Magic Numbers on page 117](#)

Benefits of PPP Fast Keepalives

- PPP fast keepalives reduce the time required for keepalive exchanges by enabling the Packet Forwarding Engine to receive LCP Echo-Request packets from the PPP subscriber and respond with LCP Echo-Reply packets, without having to send the LCP packets to the Routing Engine for processing.
- PPP fast keepalives provide increased bandwidth on the router to support a larger number of subscribers with improved performance by relieving the Routing Engine from having to process the LCP Echo-Request and Echo-Reply packets.
- PPP fast keepalives use negotiated magic numbers to identify potential traffic loopbacks to the router or network issues. You can also disable validation if needed to prevent undesired PPP session termination, for example when the PPP remote peers use arbitrary numbers rather than the negotiated number.

How PPP Fast Keepalive Processing Works

You do not need any special configuration on an MX Series router with MPCs/MICs to enable processing of PPP fast keepalive requests on the Packet Forwarding Engine. The feature is enabled by default, and cannot be disabled.

The following sequence describes how an MX Series router processes LCP Echo-Request packets and LCP Echo-Reply packets on the Packet Forwarding Engine on the MPC/MIC:

1. The Routing Engine notifies the Packet Forwarding Engine when transmission of keepalive requests is enabled on a PPP logical interface. The notification includes the magic numbers of both the server and the remote client.
2. The Packet Forwarding Engine receives the LCP Echo-Request packet initiated by the PPP subscriber (client).
3. The Packet Forwarding Engine validates the peer magic number in the LCP Echo-Request packet, and transmits the corresponding LCP Echo-Reply packet containing the magic number negotiated by the router.
4. If the Packet Forwarding Engine detects a loop condition in the link, it sends the LCP Echo-Request packet to the Routing Engine for further processing.

The Routing Engine continues to process LCP Echo-Request packets until the loop condition is cleared.

Transmission of keepalive requests from the Packet Forwarding Engine on the router is not currently enabled.

Statistics Display for PPP Fast Keepalive

When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the **Keepalive statistics** field in the output of the **show interfaces pp0.logical statistics** operational command does not include statistics for the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.

Effect of Changing the Forwarding Class Configuration

To change the default queue assignment (forwarding class) for outbound traffic generated by the Routing Engine, you can include the **forwarding-class class-name** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level.

For PPP fast (inline) keepalive LCP Echo-Request and LCP Echo-Reply packets transmitted between an MX Series router with MPCs/MICs and a PPP client, changing the forwarding class configuration takes effect immediately for both new PPP-over-Ethernet (PPPoE), PPP-over-ATM (PPPoA), and L2TP network server (LNS) subscriber sessions created after the configuration change, and for existing PPPoE, PPPoA, and LNS subscriber sessions established before the configuration change.

Preventing the Validation of PPP Magic Numbers

When the Packet Forwarding Engine validates the peer magic number in the received LCP keepalive packet, it checks whether the magic number is unexpected. The received number should match the number for the remote peer that was agreed during LCP negotiation. The remote peer number must be different than the local peer number; when they are the same, the expectation is that a loopback condition (traffic is looping back to the local peer) or some other network issue exists.

When the validation check determines that a mismatch is present, meaning that the received remote peer number is different from the negotiated number, the Packet Forwarding Engine sends the failed Echo-Reply packets to the Routing Engine. If an Echo-Reply with a valid magic number is not received within a certain interval, PPP considers this to be a keepalive failure and tears down the PPP session.

Some customer equipment might not negotiate its local magic number and instead insert an arbitrary value as the magic number it sends to the router in the keepalive packets. This number is identified as a mismatch and the session is eventually dropped. Starting in Junos OS Release 18.1R1, this result can be avoided by configuring the router to not perform a magic number validation check. Because the mismatch is never identified, the router continues to exchange PPP keepalive packets with the remote peer. To configure this behavior, include the **ignore-magic-number-mismatch** statement in an L2TP group profile, in the dynamic profile for dynamic PPP subscriber connections terminated at the router, or in the dynamic profile for dynamic tunneled PPP subscribers at the LNS.

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, this result can be avoided by configuring the router to not perform a magic number validation check.

Related Documentation

- [Configuring Keepalives](#)
- [Disabling the Sending of PPPoE Keepalive Messages](#)
- [Changing the Default Queuing and Marking of Host Outbound Traffic](#)
- [Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges on page 120](#)

Configuring Dynamic Profiles for PPP

A dynamic profile acts as a template that enables you to create, update, or remove a configuration that includes attributes for client access (for example, interface or protocol) or service (for example, IGMP). Using these profiles you can consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes simultaneously.

After they are created, the profiles reside in a profile library on the router. You can then use the **dynamic-profile** statement to attach profiles to interfaces. To assign a dynamic profile to a PPP interface, you can include the **dynamic-profile** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number ppp-options]
  dynamic-profile profile-name;
```

To monitor the configuration, issue the **show interfaces *interface-name*** command.

For information about dynamic profiles, see *Dynamic Profiles Overview* in the *Junos Subscriber Access Configuration Guide*.

For information about creating dynamic profiles, see *Configuring a Basic Dynamic Profile* in the *Junos Subscriber Access Configuration Guide*.

For information about assigning a dynamic profile to a PPP interface, see [“Attaching Dynamic Profiles to Static PPP Subscriber Interfaces”](#) on page 121 in the *Junos Subscriber Access Configuration Guide*.



NOTE: Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces for this release.

**Related
Documentation**

- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)

Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges

PPP magic numbers are negotiated between peers during LCP negotiation. The peers must have different magic numbers. When the numbers are the same, it indicates that there may be a loopback in traffic sent by the local peer. In this case, the local peer sends a new number to the remote peer. If the magic number returned by the remote peer is different than the latest number sent by the local peer, then the numbers are agreed. Otherwise, the exchange of magic numbers continues until a valid (different) number is received or the process times out, in which case the session is dropped.

After the numbers are agreed upon, the peers include their respective magic numbers when they exchange PPP keepalive (Echo-Request/Echo-Reply) packets. The Packet Forwarding Engine validates the received magic number for each exchange. A mismatch occurs when the PPP magic number received from the remote peer does not match the value agreed upon during LCP negotiation. When the validation check determines that a mismatch is present, the Packet Forwarding Engine sends the failed Echo-Request packet to the Routing Engine. If an Echo-Reply with a valid magic number is not received within a certain interval, PPP considers this to be a keepalive failure and tears down the PPP session.

In some circumstances, this behavior is not desirable. Some customer equipment does not negotiate its local magic number; instead, it inserts an arbitrary value as the magic number it sends to the router in the keepalive packets. By default, this number is identified as a mismatch and the session is eventually dropped. This result can be avoided by preventing the Packet Forwarding Engine from performing the magic number validation check. Because the mismatch is never identified, the router continues to exchange PPP keepalive packets with the remote peer.

Disable the magic number validation check by including the **ignore-magic-number-mismatch** statement as one of the PPP options applied in a dynamic PPP profile, L2TP LNS dynamic profile, or L2TP group profile. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.



NOTE: Because magic number validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends the local peer's magic number, which would indicate a loopback or other network issue. This is considered to be an unlikely situation, because LCP negotiation completed successfully, meaning no loopback was present at that time.

To configure dynamic profiles to prevent the Packet Forwarding Engine from detecting mismatches in magic numbers:

Configure the PPP option.

- For dynamic PPP subscriber connections terminated at the router:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"  
  ppp-options]
```

```
user@host# set ignore-magic-number-mismatch
```

- For dynamic tunneled PPP subscribers on LNS inline service interfaces:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit" ppp-options]
user@host# set ignore-magic-number-mismatch
```

You can use the `show ppp interface interface-name extensive` command to view whether the magic numbers are ignored.

Related Documentation

- [Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests on page 116](#)
- For L2TP group profiles: [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

You can attach a dynamic profile to a static PPP subscriber interface. When a PPP subscriber logs in, the specified dynamic profile is instantiated and the services defined in the profile are applied to the interface.

To attach a dynamic profile to a static PPP subscriber interface:

1. Specify that you want to configure PPP options.

```
[edit interfaces pp0 unit 0]
user@host# edit ppp-options
```

2. Specify the dynamic profile you want to associate with the interface.

```
[edit interfaces pp0 unit 0 ppp-options]
user@host# set dynamic-profile vod-profile-50
```

Related Documentation

- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Dynamic Profiles Overview](#)
- [Configuring a Basic Dynamic Profile](#)
- [Example: Minimum PPPoE Dynamic Profile on page 121](#)
- [Verifying and Managing PPP Configuration for Subscriber Management on page 143](#)

Example: Minimum PPPoE Dynamic Profile

This example shows the minimum configuration for a dynamic profile that is used for static PPPoE interfaces. The configuration must include the `interfaces pp0` stanza.

```
dynamic-profiles {
  ppp-profile-1 {
    interfaces {
```

```
pp0 {  
  unit "$junos-interface-unit";  
}  
}  
}
```

**Related
Documentation**

- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121](#)

Applying RADIUS Route Attributes to Subscribers or Access Networks

- [Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 123](#)
- [Verifying the Configuration of Access and Access-Internal Routes for PPP Subscribers on page 124](#)

Configuring Dynamic Access-Internal Routes for PPP Subscriber Management

You can dynamically configure access-internal routes for PPP subscribers. Configuring support for access-internal variables is optional, but it ensures that the values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

Starting in Junos OS Release 15.1R1, we no longer recommend that you always include the **access-internal** stanza in the dynamic-profile when the **access** stanza is present for framed route support. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute (22) or Framed-IPv6-Route attribute [99].

For PPP subscriber interfaces, you do not need to specify the MAC address for access-internal routes.

To dynamically configure access-internal routes for PPP:

1. Specify that you want to configure the access-internal route.

```
user@host# edit dynamic-profiles profile-name routing-options
```

2. Specify the IP address as a variable.

```
[edit dynamic-profiles profile-name routing-options]  
user@host# edit access-internal route $junos-subscriber-ip-address
```

3. Specify the qualified-next-hop as a variable.

```
[edit dynamic-profiles profile-name routing-options access-internal route  
$junos-subscriber-ip-address]
```

```
user@host# set qualified-next-hop $junos-interface-name
```

Release History Table

Release	Description
15.1R1	Starting in Junos OS Release 15.1R1, we no longer recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

Related Documentation

- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)
- [Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 111](#)

Verifying the Configuration of Access and Access-Internal Routes for PPP Subscribers

Purpose View configuration information for access routes and access-internal routes on PPP subscribers. The access-internal routes are those that are automatically installed when a client profile is instantiated.

Action • To display extensive information about access routes and access-internal routes:

```
user@host>show route extensive
```

- To display the configuration for access routes:

```
user@host>show route protocol access
```

- To display the configuration for access-internal routes:

```
user@host> show route protocol access-internal
```

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

Configuring Authentication for PPP

- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Modifying the CHAP Challenge Length on page 127](#)

Configuring Dynamic Authentication for PPP Subscribers

You can configure a dynamic profile that includes PPP authentication that enables PPP clients to dynamically access the network. You can specify either CHAP or PAP authentication. Optionally, you can also control the order in which the router negotiates the CHAP and PAP protocols.

For dynamic interfaces, the router supports unidirectional authentication only—the router always functions as the authenticator. When you configure PPP authentication in a dynamic profile, CHAP authentication supports the **challenge-length** option, which enables you to configure the minimum length and maximum length of the CHAP challenge message. Neither CHAP authentication nor PAP authentication supports any other configuration options, including the **passive** statement.



NOTE: Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces.

To configure authentication in a dynamic profile for PPP subscriber interfaces:

1. Name the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles vod-profile-25
```

2. Configure the interfaces and unit for the dynamic profile. Use **pp0** for the interface type and the predefined variable for the unit.

```
[edit dynamic-profiles vod-profile-25]
user@host# edit interfaces pp0 unit $junos-interface-unit
```

3. Configure PPP options.

```
[edit dynamic-profiles vod-profile-25 interfaces pp0 unit "$junos-interface-unit"]
user@host# edit ppp-options
```

4. Specify the authentication protocol used in the dynamic profile. You can configure either CHAP or PAP. There are no additional options for either authentication protocol.

```
[edit dynamic-profiles vod-profile-25 interfaces pp0 unit "$junos-interface-unit"  
  ppp-options]  
user@host# set chap
```

5. (Optional) Configure the minimum length and maximum length of the CHAP challenge message.

See [“Modifying the CHAP Challenge Length” on page 127](#).

6. (Optional) Configure the order in which the router negotiates the CHAP and PAP authentication protocols.

See [“Controlling the Negotiation Order of PPP Authentication Protocols” on page 132](#).

7. (Optional) Configure the router to prompt the CPE to negotiate the DNS addresses for dynamic PPPoE subscribers during IPCP negotiation.

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"  
  ppp-options]  
user@host# set ipcp-suggest-dns-option
```

**Related
Documentation**

- [Modifying the CHAP Challenge Length on page 127](#)
- [Controlling the Negotiation Order of PPP Authentication Protocols on page 132](#)
- [Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 136](#)
- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- [Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121](#)
- [Dynamic Profiles Overview](#)
- [Configuring a Basic Dynamic Profile](#)
- [Example: Minimum PPPoE Dynamic Profile on page 121](#)
- [Verifying and Managing PPP Configuration for Subscriber Management on page 143](#)

Modifying the CHAP Challenge Length

You can modify the default minimum length and maximum length of the Challenge Handshake Authentication Protocol (CHAP) challenge message that the router sends to a PPP client. The CHAP challenge message, which contains information that is unique to a particular PPP subscriber session, is used as part of the authentication mechanism between the router and the client to verify the identity of the client for access to the router.

By default, the minimum length of the CHAP challenge is 16 bytes, and the maximum length is 32 bytes. You can override this default to configure the CHAP challenge minimum length and maximum length in the range 8 bytes through 63 bytes.



BEST PRACTICE: We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.

Before you begin:

- Configure the CHAP protocol on the interface.
 - For dynamic PPP subscriber interfaces, see [“Configuring Dynamic Authentication for PPP Subscribers” on page 125](#).
 - For static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

To configure the minimum and maximum length of the CHAP challenge message:

1. Specify that you want to configure PPP options.

- For dynamic PPP subscriber interfaces:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]
user@host# edit ppp-options
```

- For static interfaces with PPP encapsulation:

```
[edit interfaces pp0 unit logical-unit-number]
user@host# edit ppp-options
```

2. Specify that you want to configure CHAP options.

- For dynamic PPP subscriber interfaces:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"
  ppp-options]
user@host# edit chap
```

- For static interfaces with PPP encapsulation:

```
[edit interfaces pp0 unit logical-unit-number ppp-options]
user@host# edit chap
```

3. Specify the minimum length and maximum length of the CHAP challenge.

- For dynamic PPP subscriber interfaces:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"  
  ppp-options chap]  
user@host# set challenge-length minimum minimum-length maximum  
  maximum-length
```

- For static interfaces with PPP encapsulation:

```
[edit interfaces pp0 unit logical-unit-number ppp-options chap]  
user@host# set challenge-length minimum minimum-length maximum  
  maximum-length
```

For example, the following **challenge-length** statement in a dynamic profile named `pppoe-client-profile` sets the minimum length of the CHAP challenge to 20 bytes, and the maximum length to 40 bytes.

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"  
  ppp-options chap]  
user@host# set challenge-length minimum 20 maximum 40
```

**Related
Documentation**

- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- [Configuring the PPP Challenge Handshake Authentication Protocol](#)

Configuring PPP Network Control Protocol Negotiation

- [PPP Network Control Protocol Negotiation Mode Overview on page 129](#)
- [Controlling the Negotiation Order of PPP Authentication Protocols on page 132](#)
- [Configuring the PPP Network Control Protocol Negotiation Mode on page 134](#)
- [Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 136](#)

PPP Network Control Protocol Negotiation Mode Overview

The *Network Control Protocol* (NCP) is a mechanism used to establish and configure different Network Layer protocols for Point-to-Point Protocol (PPP) connections. Starting in Junos OS Release 14.1, on MX Series routers with Modular Port Concentrators (MPCs), you can configure *PPP NCP negotiation* to actively or passively control subscriber connections initiated by the router functioning as a PPP server.

- [PPP NCP Negotiation Modes on page 129](#)
- [PPP NCP Negotiation Mode Supported Configurations on page 130](#)
- [PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers on page 130](#)
- [PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers on page 131](#)
- [PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations on page 131](#)

PPP NCP Negotiation Modes

PPP NCP negotiation operates in either of the following modes:

- *Active PPP NCP negotiation mode*—The router sends an NCP Configuration Request message without waiting for the PPP client to do so.
- *Passive PPP NCP negotiation mode*—The router waits for the PPP client to send an NCP Configuration Request message before sending its own Configuration Request message. Dynamic subscriber interface connections and static subscriber interface connections use passive PPP NCP negotiation by default.

Router behavior for active mode and passive mode PPP NCP negotiation differs for dynamic PPP subscribers and static PPP subscribers, as summarized in [Table 8 on page 130](#).

Table 8: PPP NCP Negotiation Mode Behavior for Dynamic and Static Subscribers

PPP Subscribers	PPP NCP Negotiation Mode	Router Behavior
Dynamic	Active	The router establishes the local network address and uses it to send the NCP Configuration Request message without waiting for the PPP client to send a Configuration Request.
Dynamic	Passive	The router establishes the local network address after it receives the NCP Configuration Request message from the PPP client.
Static	Active	The router sends the authentication acknowledgement to the PPP client, and then sends the NCP Configuration Request message without waiting for the PPP client to send its own Configuration Request.
Static	Passive	The router sends the authentication acknowledgement to the PPP client, and then waits for an NCP Configuration Request message from the client before sending a Configuration Request.

PPP NCP Negotiation Mode Supported Configurations

You can configure PPP Network Control Protocol (NCP) negotiation for the following single-stack and dual-stack subscriber configurations on MX Series routers with MPCs:

- Dynamic PPP subscriber connections terminated at the router
- Static PPP subscriber connections terminated at the router
- Dynamic tunneled PPP subscribers at the L2TP network server (LNS)
- Static tunneled PPP subscribers at the L2TP network server (LNS) on an inline service (si) interface

PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers

To configure active PPP IPv4 Network Control Protocol (IPNCP) negotiation for dynamic and static PPP subscribers in a single-stack or dual-stack configuration, make sure you meet the following requirements:

- Configure the IPv4 (**inet**) protocol family in a dynamic profile (for dynamic subscribers) or at the interface level (for static subscribers).
- Assign any of the following IPv4 address attributes for the subscriber during the authentication process:

- Framed-IP-Address (RADIUS Attribute 8)—RADIUS explicit IPv4 address
- Framed-Pool (RADIUS Attribute 88)—RADIUS IPv4 address pool name
- IPv4 attributes allocated from a locally configured address pool

When you have met these requirements, use the **initiate-ncp ip** statement to enable active IPNCP negotiation for dynamic and static subscribers in a single-stack or dual-stack configuration.

PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers

To configure active PPP IPv6 Network Control Protocol (IPv6NCP) negotiation for dynamic and static PPP subscribers in a single-stack or dual-stack configuration, make sure you meet the following requirements:

- Configure the IPv6 (**inet6**) protocol family in a dynamic profile (for dynamic subscribers) or at the interface level (for static subscriber).
- Assign any of the following IPv6 address attributes for the subscriber during the authentication process:
 - Delegated-IPv6-Prefix (RADIUS Attribute 123)—RADIUS explicit IPv6 address
 - Framed-IPv6-Prefix (RADIUS Attribute 97)—RADIUS explicit IPv6 prefix
 - Framed-IPv6-Pool (RADIUS Attribute 100)—RADIUS explicit IPv6 address or prefix pool name
 - IPv6 attributes allocated from a locally configured Neighbor Discovery Router Advertisement (NDRA) pool

When you have met these requirements, use the **initiate-ncp ipv6** statement to enable active IPv6NCP negotiation for dynamic and static subscribers in a single-stack or dual-stack configuration.

PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations

You can configure either active or passive PPP NCP negotiation for the IPv4 and IPv6 subscriber interfaces in a dual-stack configuration.

To configure active negotiation in a dual-stack configuration, do all of the following:

- Make sure you meet the IPv4 and IPv6 protocol and address family requirements.
- Use the **initiate-ncp ip** statement to enable active negotiation for the IPv4 subscriber interface.
- Use the **initiate-ncp ipv6** statement to enable active negotiation for the IPv6 subscriber interface.

To configure passive negotiation in a dual-stack configuration, do both of the following:

- Make sure you meet the IPv4 and IPv6 protocol and address family requirements.

- Use the **initiate-ncp dual-stack-passive** statement to enable passive negotiation for the dual-stack configuration. The **initiate-ncp dual-stack-passive** statement overrides the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured.

The following additional guidelines apply when you configure PPP NCP negotiation for dual-stack subscribers:

- Dual-stack subscribers configured for either active mode or passive mode PPP NCP negotiation continue to use the same negotiation mode when the NCP mechanism is renegotiated.
- Using the **on-demand-ip-address** statement to save IPv4 addresses for dual-stack PPP subscribers when you are not using the IPv4 service has no effect on configuration of the PPP NCP negotiation mode in a dual-stack configuration.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, on MX Series routers with Modular Port Concentrators (MPCs), you can configure <i>PPP NCP negotiation</i> to actively or passively control subscriber connections initiated by the router functioning as a PPP server.

Related Documentation

- [Configuring the PPP Network Control Protocol Negotiation Mode on page 134](#)

Controlling the Negotiation Order of PPP Authentication Protocols

You can control the order in which the router tries to negotiate PPP authentication protocols when it verifies that a PPP client can access the network. By default, the router first tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication. If the attempt to negotiate CHAP authentication is unsuccessful, the router then tries to negotiate Password Authentication Protocol (PAP) authentication.

You can modify this default negotiation order in any of the following ways:

- Specify that the router negotiate PAP authentication first, followed by CHAP authentication if PAP negotiation is unsuccessful.

When you specify both authentication protocols in either order, you must enclose the set of protocol names in square brackets ([]).

- Specify that the router negotiate only CHAP authentication.
- Specify that the router negotiate only PAP authentication.

Before you begin:

- Configure the CHAP or PAP protocol on the interface.
 - For dynamic PPP subscriber interfaces, see [“Configuring Dynamic Authentication for PPP Subscribers” on page 125](#).

- For CHAP on static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.
- For PAP on static interfaces with PPP encapsulation, see *Configuring the PPP Password Authentication Protocol On a Physical Interface*.

To control the order in which the router negotiates PPP authentication protocols:

1. Specify that you want to configure PPP options.

- For dynamic PPP subscriber interfaces:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]
user@host# edit ppp-options
```

- For static interfaces with PPP encapsulation:

```
[edit interfaces pp0 unit logical-unit-number]
user@host# edit ppp-options
```

2. Specify the negotiation order for PPP authentication protocols on the router.

- For dynamic PPP subscriber interfaces:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"
ppp-options]
user@host# set authentication [authentication-protocols]
```

- For static interfaces with PPP encapsulation:

```
[edit interfaces pp0 unit logical-unit-number ppp-options]
user@host# set authentication [authentication-protocols]
```

The following sample **authentication** statements in a dynamic profile named pppoe-client-profile show the different ways you can configure the negotiation order for PPP authentication protocols. (The **authentication** statements for configuring static interfaces are identical.)

- To specify that the router negotiate PAP authentication first, followed by CHAP authentication:

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
ppp-options]
user@host# set authentication [pap chap]
```

- To specify that the router negotiate only CHAP authentication:

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
ppp-options]
user@host# set authentication chap
```

- To specify that the router negotiate only PAP authentication:

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
ppp-options]
user@host# set authentication pap
```

- To restore the default negotiation order for PPP authentication protocols after you have modified it:

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"  
  ppp-options]  
user@host# set authentication [chap pap]
```

**Related
Documentation**

- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol On a Physical Interface*

Configuring the PPP Network Control Protocol Negotiation Mode

Starting in Junos OS Release 14.1, configuring PPP Network Control Protocol (NCP) negotiation enables you to actively or passively control subscriber connections initiated by the router functioning as a PPP server. Both dynamic and static subscriber interface connections use passive PPP NCP negotiation by default.

You can configure the PPP NCP negotiation mode (active or passive) for the following subscriber configurations on MX Series routers with MPCs:

- Dynamic PPP subscriber connections terminated at the router, using a dynamic profile
- Static PPP subscriber connections terminated at the router, using a per-interface configuration
- Dynamic tunneled PPP subscribers at the L2TP network server (LNS), using a dynamic profile
- Static tunneled PPP subscribers at the LNS, using a per-inline service (si) interface configuration
- Dynamic and static tunneled PPP subscribers at the LNS, using a user-group profile

To configure PPP NCP negotiation mode:

1. Specify that you want to configure PPP-specific properties for the subscriber.

- For dynamic PPP subscriber connections terminated at the router:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]  
user@host# edit ppp-options
```

- For static PPP subscriber connections terminated at the router:

```
[edit interfaces pp0 unit logical-unit-number]  
user@host# edit ppp-options
```

- For dynamic tunneled PPP subscribers at the LNS:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
  "$junos-interface-unit"]  
user@host# edit ppp-options
```

- For static tunneled PPP subscribers at the LNS:

```
[edit interfaces si-fpc/pic/port unit logical-unit-number]
user@host# edit ppp-options
```

- In a group profile for dynamic and static tunneled PPP subscribers at the LNS:

```
[edit access group-profile profile-name ppp]
user@host# edit ppp-options
```

2. Configure PPP NCP negotiation mode in any of the following ways:

- To configure active PPP NCP negotiation for IPv4 subscribers in a single-stack or dual-stack configuration, use the **initiate-ncp ip** statement.

For example, to configure active negotiation for static IPv4 connections terminated at the router:

```
[edit interfaces pp0 unit logical-unit-number ppp-options]
user@host# initiate-ncp ip
```

- To configure active PPP NCP negotiation for IPv6 subscribers in a single-stack or dual-stack configuration, use the **initiate-ncp ipv6** statement.

For example, to configure active negotiation for dynamic IPv6 connections terminated at the router:

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"
 ppp-options]
user@host# initiate-ncp ipv6
```

- To configure passive PPP NCP negotiation for dynamic or static subscribers in an IPv4 and IPv6 dual-stack configuration, use the **initiate-ncp dual-stack-passive** statement, which overrides both the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured.

For example, to configure passive negotiation for dynamic tunneled PPP subscribers at the LNS in an IPv4 and IPv6 dual-stack configuration:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
 "$junos-interface-unit"]
user@host# initiate-ncp dual-stack-passive
```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, configuring PPP Network Control Protocol (NCP) negotiation enables you to actively or passively control subscriber connections initiated by the router functioning as a PPP server.

Related Documentation

- [PPP Network Control Protocol Negotiation Mode Overview on page 129](#)

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses

Starting in Junos OS Release 15.1, you can configure a router to prompt any customer premises equipment (CPE) to send the IPv4 primary or secondary DNS address options in the next configuration request if the options are not included in an initial IPCP configuration request during IPCP negotiations or if the router rejects the request. This DNS option enables the router to control IPv4 DNS address provisioning for dynamic and static, terminated PPPoE and LNS subscribers. The router includes the address options in the IPCP configuration NAK message that it sends to the CPE. The CPE then negotiates both primary and secondary IPv4 DNS addresses. Using this option avoids a situation in which the CPE does not take advantage of the DNS addresses available at the router.

To configure the router to prompt the CPE to negotiate the DNS addresses for dynamic PPPoE subscribers:

- Specify the DNS negotiation option.

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"  
  ppp-options]  
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for static PPPoE subscribers:

- Specify the DNS negotiation option.

```
[edit interfaces interface-name ppp-options]  
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for dynamic LNS subscribers:

- Specify the DNS negotiation option.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
  "$junos-interface-unit" ppp-options]  
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for static LNS subscribers:

- Specify the DNS negotiation option.

```
[edit interfaces si-slot/pic/port unit logical-unit-number ppp-options]  
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for tunneled PPP subscribers with an LNS user group profile:

- Specify the DNS negotiation option.

```
[edit access group-profile profile-name ppp-options]  
user@host# set ipcp-suggest-dns-option
```

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, you can configure a router to prompt any customer premises equipment (CPE) to send the IPv4 primary or secondary DNS address options in the next configuration request if the options are not included in an initial IPCP configuration request during IPCP negotiations or if the router rejects the request.

Related Documentation

- [Configuring the Group Profile for Defining L2TP Attributes](#)
- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)
- [Dynamic Profiles Overview](#)
- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)

Configuring High Availability in the PPP Access Network

- [High Availability Using Unified ISSU in the PPP Access Network on page 139](#)
- [Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 140](#)
- [Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 141](#)

High Availability Using Unified ISSU in the PPP Access Network

The unified in-service software upgrade (unified ISSU) feature supports the PPPoE access model used by subscriber management. This support ensures that the router preserves active PPPoE subscriber sessions and session services after a unified ISSU has completed.

See *Getting Started with Unified In-Service Software Upgrade* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

Unified ISSU supports the subscriber management PPPoE access model for static and dynamic PPPoE access, and includes the following features:

- Terminated, non-tunneled PPPoE connections configured with static or dynamic PPP logical interfaces and static or dynamic underlying interfaces
- Subscriber services on single-link PPP interfaces
- Preservation of statistics for accounting, filter, and CoS on MPC/MIC interfaces

Unified ISSU for the subscriber management PPPoE access model *does not support* Multilink Point-to-Point Protocol (MLPPP) bundle interfaces. MLPPP bundle interfaces require the use of an Adaptive Services PIC or Multiservices PIC to provide PPP subscriber services. These PICs do not support unified ISSU.

- Related Documentation**
- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 111](#)
 - [Unified ISSU System Requirements](#)

Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to wait 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover (GRES) takes place.

- [Benefits of Delaying Removal of Access Routes and Access-Internal Routes on page 140](#)
- [Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes on page 140](#)
- [Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes on page 140](#)

Benefits of Delaying Removal of Access Routes and Access-Internal Routes

The 3-minute delay in removing access routes and access-internal routes after a graceful Routing Engine switchover provides sufficient time for the DHCP client process (jdhcpd), PPP client process (jpppd), or routing protocol process (rpd) to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

Configuring the router to delay removal of access routes and access-internal routes after a graceful Routing Engine switchover has the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine
- Prevents loss of subscriber traffic due to unavailable routes

Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale access routes and access-internal routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the jdhcpd or jpppd client process to reinstall all of the subscriber routes.

Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process (rpd) immediately purges the stale

access routes and access-internal routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

**Related
Documentation**

- [Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 108](#)
- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover

In subscriber networks configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to delay for 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover takes place.

To configure the router to delay removal (flushing) of access-routes and access-internal routes after a graceful Routing Engine switchover:

1. Specify that you want to configure subscriber management.

```
[edit system services]  
user@host# edit subscriber-management
```

2. Configure the router to wait 180 seconds before removing access-routes and access-internal routes after a graceful Routing Engine switchover.

```
[edit system services subscriber-management]  
user@host# set gres-route-flush-delay
```

**Related
Documentation**

- [Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 107](#)
- [Access and Access-Internal Routes for Subscriber Management on page 31](#)
- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

CHAPTER 16

Monitoring and Managing PPP for Subscriber Access

- [Verifying and Managing PPP Configuration for Subscriber Management on page 143](#)
- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 143](#)

Verifying and Managing PPP Configuration for Subscriber Management

Purpose View or clear information about PPP configuration for subscriber management.

Action • To display information about PPP interfaces:

user@host> [show ppp interface](#)

• To display PPP statistics information:

user@host> [show ppp statistics](#)

• To display PPP session summary information:

user@host> [show ppp summary](#)

• To display PPP address-pool information:

user@host> [show ppp address-pool](#)

Related Documentation • [Dynamic Profiles for PPP Subscriber Interfaces Overview on page 115](#)
• [CLI Explorer](#)

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane quiescing has not started). The second example shows that unified ISSU

is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
```

General:

Graceful Restart	Enabled
Mastership	Master
Database	Available
Chassisd ISSU State	IDLE
ISSU State	IDLE
ISSU Wait	0

```
user@host> show system subscriber-management summary
```

General:

Graceful Restart	Enabled
Mastership	Master
Database	Available
Chassisd ISSU State	DAEMON_ISSU_PREPARE
ISSU State	PREPARE
ISSU Wait	198

**Related
Documentation**

- [High Availability Using Unified ISSU in the PPP Access Network on page 139](#)
- [High Availability Using Unified ISSU in the DHCP Access Network on page 106](#)
- [High Availability Using Unified ISSU in the L2TP Access Network on page 297](#)
- *Getting Started with Unified In-Service Software Upgrade*

PART 4

Configuring the L2TP Access Network

- [L2TP and Subscriber Access Overview on page 147](#)
- [Configuring L2TP Tunneling and Switching for Subscribers on page 153](#)
- [Configuring L2TP Control Messages for Subscribers on page 169](#)
- [Configuring L2TP LAC Subscribers on page 173](#)
- [Configuring L2TP LAC Tunneling for Subscribers on page 183](#)
- [Configuring Use of Subscriber Access Line and Connect Speed Information on page 213](#)
- [Configuring L2TP LNS Inline Service Interfaces on page 239](#)
- [Configuring IP Packet Fragment Reassembly on page 289](#)
- [Configuring High Availability in the L2TP Access Network on page 293](#)
- [Monitoring and Managing L2TP for Subscriber Access on page 299](#)

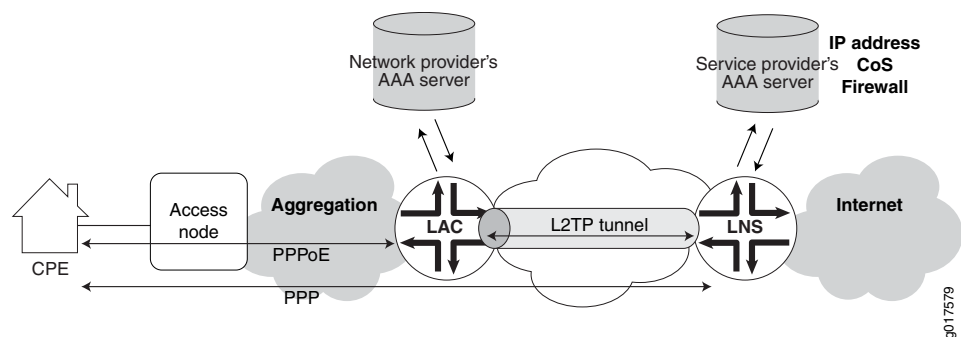
L2TP and Subscriber Access Overview

- [L2TP for Subscriber Access Overview on page 147](#)
- [L2TP Terminology on page 149](#)
- [L2TP Implementation on page 150](#)

L2TP for Subscriber Access Overview

The Layer 2 Tunneling Protocol (L2TP) is a client-server protocol that allows the Point-to-Point Protocol (PPP) to be tunneled across a network. L2TP encapsulates Layer 2 packets, such as PPP, for transmission across a network. An L2TP access concentrator (LAC), configured on an access device, receives packets from a remote client and forwards them to an L2TP network server (LNS) on a remote network. The LNS functions as the logical termination point of the PPP session tunneled by the LAC from the remote client. [Figure 8 on page 147](#) shows a simple L2TP topology.

Figure 8: Typical L2TP Topology




L2TP separates the termination of access technologies, such as cable or xDSL, from the termination of PPP and subsequent access to a network. This separation enables public ISPs to outsource their access technologies to competitive local exchange carriers (CLECs). L2TP provides ISPs the capability to supply VPN service; private enterprises can reduce or avoid investment in access technologies for remote workers.

You can configure your router to act as the LAC in PPP pass-through mode in which the LAC receives packets from a remote client and then forwards them at Layer 2 directly to the LNS. The PPP session is terminated on the LNS. This LAC implementation supports only Point-to-Point Protocol over Ethernet (PPPoE) subscribers over dynamic or static

logical interfaces. [Figure 9 on page 148](#) shows the protocol layer stacking for an L2TP pass-through connection.

Figure 9: Protocol Stacking for L2TP Subscribers in Pass-Through Mode



NOTE: On MX Series routers, the LAC and 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](#)

Certain M Series routers support LNS functions on services PICs. For more information about the L2TP implementation on M Series routers, see [L2TP Services Configuration Overview](#).

The LAC dynamically creates tunnels based on AAA authentication parameters and transmits L2TP packets to the LNS by means of the IP/User Datagram Protocol (UDP). Traffic travels in an L2TP *session*; a tunnel is an aggregation of one or more sessions. You can also provision a domain map that is used by AAA to determine whether to tunnel or terminate the PPPoE subscriber on the LAC. A one-to-one mapping exists between each PPP subscriber tunneled to the LNS and an L2TP session.

When the LNS is an MX Series router, a LAC-facing peer interface on an MPC provides an IP address for the exchange of IP packets between the tunnel endpoints; the Routing Engine maintains the L2TP tunnels. The Packet Forwarding Engine hosts one or more inline services (si) interfaces. These interfaces function like a virtual physical interface and *anchor* the L2TP sessions on the LNS. The si interface enables L2TP services without requiring a special services PIC. Finally, another interface is used to transmit the subscriber data to and from the Internet.

The characteristics of the tunnel can originate either from a tunnel profile that you configure or from RADIUS tunnel attributes and vendor-specific attributes (VSAs) from the AAA server accessible at the LAC. You can include a tunnel profile in a domain map, which applies the tunnel profile before RADIUS authentication takes place. You can use RADIUS standard attributes and VSAs to override any or all characteristics configured by the tunnel profile in a domain map. Alternatively, RADIUS can itself apply a tunnel profile when the RADIUS Tunnel-Group VSA [26-64] is specified in the RADIUS login.

The Virtual-Router VSA [26-1] in the subscriber profile on the service provider AAA server (accessible from the LNS) determines the routing instance in which the L2TP session is brought up on the LNS. When this VSA is not present, the subscriber session comes up in the same routing instance as the tunnel, because the AAA server can be accessed only from the routing instance in which the tunnel terminates on the LNS.

This behavior is different than for DHCP and non-tunneled PPPoE subscribers, which come up in the default routing instance in the absence of the Virtual-Router VSA. For L2TP subscribers, you must include this VSA in the subscriber profile when you want the subscriber session to come up in a different routing instance than the tunnel routing instance.

Starting in Junos OS Release 17.4R1, The LNS includes the following RADIUS attributes when it sends an Access-Request message to the RADIUS server:

- Tunnel-Type (64)
- Tunnel-Medium-Type (65)
- Tunnel-Client-Endpoint (66)
- Tunnel-Server-Endpoint (67)
- Acct-Tunnel-Connection (68)
- Tunnel-Assignment-Id (82)
- Tunnel-Client-Auth-Id (90)
- Tunnel-Server-Auth-Id (91)

In earlier releases, the LNS includes those attributes only in the accounting records it sends to the RADIUS server. In the Access-Request messages, they can be used to correlate on the RADIUS server the session from the LAC to the LNS.

The LAC supports RADIUS-initiated mirroring, which creates secure policies based on certain RADIUS VSAs, and uses RADIUS attributes to identify a subscriber whose traffic is to be mirrored. (This feature is not supported for an LNS configured on an MX Series router.)

The LAC and the LNS support unified ISSU. When an upgrade is initiated, the LAC completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade. Subscriber logouts are recorded during the upgrade and are completed after the upgrade has completed.

Related Documentation

- *RADIUS IETF Attributes Supported by the AAA Service Framework*
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)
- *Domain Mapping Overview*
- *Getting Started with Unified In-Service Software Upgrade*

L2TP Terminology

[Table 9 on page 150](#) describes the basic terms for L2TP.

Table 9: L2TP Terms

Term	Description
AVP	Attribute value pair (AVP)—Combination of a unique attribute—represented by an integer—and a value containing the actual value identified by the attribute.
Call	A connection (or attempted connection) between a remote system and the LAC.
LAC	L2TP access concentrator (LAC)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LNS. The LAC sits between an LNS and a remote system and forwards packets to and from each.
LNS	L2TP network server (LNS)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. The LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC.
Peer	In the L2TP context, either the LAC or LNS. The LAC's peer is an LNS, and vice versa.
Proxy authentication	PPP pre-authentication performed by the LAC on behalf of the LNS. The proxy data is sent by the LAC to the LNS containing attributes such as authentication type, authentication name, and authentication challenge. The LNS responds with the authentication results.
Proxy LCP	Link Control Protocol (LCP) negotiation that is performed by the LAC on behalf of the LNS. The proxy is sent by the LAC to the LNS containing attributes such as the last configuration attributes sent and received from the client.
Remote system	An end system or router attached to a remote access network, which is either the initiator or recipient of a call.
Session	A logical connection created between the LAC and the LNS when an end-to-end PPP connection is established between a remote system and the LNS. NOTE: There is a one-to-one relationship between established L2TP sessions and their associated PPP connections.
Tunnel	A connection between the LAC-LNS pair consisting of a control connection and 0 or more L2TP sessions.

Related Documentation

- [L2TP for Subscriber Access Overview on page 147](#)

L2TP Implementation

L2TP is implemented on four levels:

- Source—The local router acting as the LAC.
- Destination—The remote router acting as the LNS.

- Tunnel—A direct path between the LAC and the LNS.
- Session—A PPP connection in a tunnel.

When the router has established destinations, tunnels, and sessions, you can control the L2TP traffic. Making a change to a destination affects all tunnels and sessions to that destination; making a change to a tunnel affects all sessions in that tunnel. For example, closing a destination closes all tunnels and sessions to that destination.

Sequence of Events on the LAC

The router acting as the LAC creates destinations, tunnels, and sessions dynamically, as follows:

1. The client initiates a PPP connection with the router.
2. The router and the client exchange Link Control Protocol (LCP) packets. The LAC negotiates on behalf of the LNS; this is known as *proxy LCP*.
3. The LAC authenticates the client on behalf of the LNS; this is known as *proxy authentication*. By using either a local database related to the domain name or RADIUS authentication, the router determines either to terminate or to tunnel the PPP connection.
4. If the router discovers that it should tunnel the session, it does the following:
 - a. Sets up a new destination or selects an existing destination.
 - b. Sets up a new tunnel or selects an existing tunnel.

When a shared secret is configured in either the tunnel profile or the RADIUS attribute Tunnel-Password [69]—depending on which method is used to configure the tunnel—the secret is used to authenticate the tunnel during the establishment phase. The LAC includes the Challenge AVP in the SCCRQ message sent to the LNS. The LNS returns the Challenge Response AVP in the SCCRQ message. If the response from the LNS does not match the value expected by the LAC, then tunnel authentication fails and the tunnel is not established.

- c. Opens a new session.
5. The router forwards the results of the LCP negotiations and authentication to the LNS.

A PPP connection now exists between the client and the LNS.



NOTE: The router discards received packets if the size of the variable-length, optional offset pad field in the L2TP header is too large. The router always supports packets that have an offset pad field of up to 16 bytes, and may support larger offset pad fields, depending on other information in the header. This restriction is a possible, although unlikely, cause of excessive discarding of L2TP packets.



NOTE: When the LAC terminates a PPP session, it generates a PPP disconnect cause and includes this information in the PPP Disconnect Cause Code (AVP 46) when it sends a Call-Disconnect-Notify (CDN) message to the LNS. The code value is 0, which indicates a global error with no information available.

Sequence of Events on the LNS

A router acting as an LNS might be set up as follows:

1. The LAC initiates a tunnel with the router acting as the LNS.
2. The LNS verifies that a tunnel with this LAC is valid: the destination is configured, the hostname and the tunnel password are correct.
3. The LNS completes the tunnel setup with the LAC.
4. The LAC sets up a session and initiates a session request to the LNS.
5. The LNS uses a static interface or creates a dynamic interface to anchor the PPP session.
6. If they are enabled and present, the LNS accepts the proxy LCP and the proxy authentication data and passes them to PPP.
7. PPP processes the proxy LCP, if it is present, and, if the proxy LCP is acceptable, places LCP on the LNS in opened state without renegotiation of LCP.
8. PPP processes the proxy authentication data, if it is present, and passes the data to AAA for verification. (If the data is not present, PPP requests the data from the peer.)



NOTE: When the proxy LCP is not present or not acceptable, the LNS negotiates LCP with the peer. When LCP renegotiation is enabled on the LNS, the LNS ignores any pre-negotiated LCP parameters and renegotiates both the LCP parameters and PPP authentication with the PPP client.

-
9. The LNS passes the authentication results to the peer.

Related Documentation

- [L2TP for Subscriber Access Overview on page 147](#)

CHAPTER 18

Configuring L2TP Tunneling and Switching for Subscribers

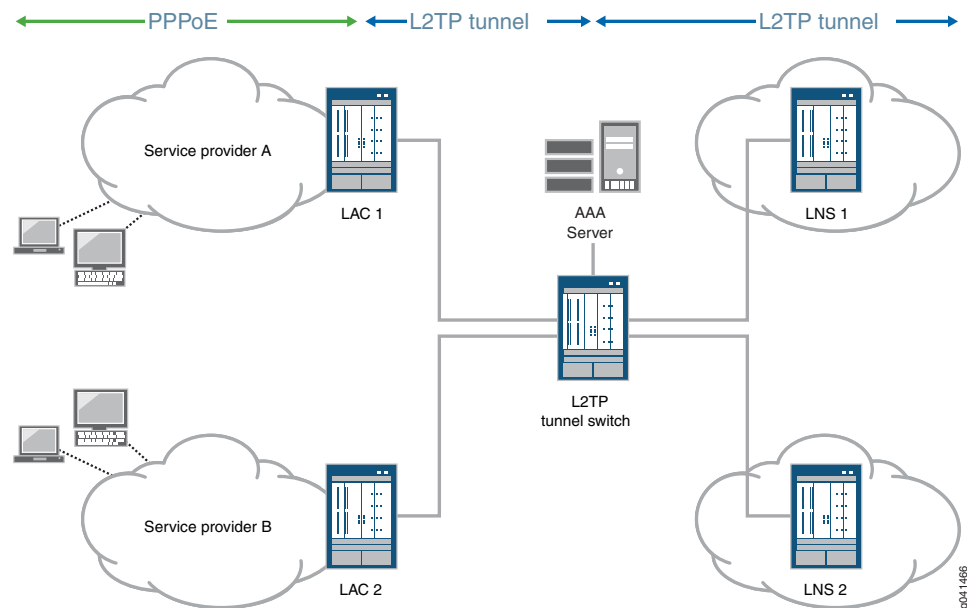
- [L2TP Tunnel Switching Overview on page 153](#)
- [Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 157](#)
- [Configuring L2TP Tunnel Switching on page 161](#)
- [Setting the L2TP Receive Window Size on page 163](#)
- [Setting the L2TP Tunnel Idle Timeout on page 163](#)
- [Setting the L2TP Destruct Timeout on page 164](#)
- [Configuring the L2TP Destination Lockout Timeout on page 165](#)
- [Removing an L2TP Destination from the Destination Lockout List on page 165](#)
- [Configuring L2TP Drain on page 166](#)
- [Using the Same L2TP Tunnel for Injection and Duplication of IP Packets on page 167](#)

L2TP Tunnel Switching Overview

L2TP tunnel switching, also known as L2TP multihop, simplifies the deployment of an L2TP network across multiple domains. A router that lies between a LAC and an LNS is configured as an *L2TP tunnel switch* (LTS)—sometimes referred to simply as a *tunnel switch* or a *tunnel switching aggregator* (TSA)—as shown in [Figure 10 on page 154](#). The LTS is configured as both an LNS and a LAC. When a remote LAC sends encapsulated PPP packets to the LNS configured on the LTS, the LTS can forward or redirect the packets through a different tunnel to a different LNS beyond the LTS. The logical termination point of the original L2TP session is switched to a different endpoint.

For example, in the network shown in [Figure 10 on page 154](#), packets from the subscriber provisioned by service provider A are initially targeted at the LNS configured on the LTS. The LTS might redirect those packets to LNS1.

Figure 10: L2TP Tunnel Switching Network Topology



L2TP tunnel switching simplifies network configuration when the administrative domain of a LAC is different from that of the desired LNS. For example:

- The LTS acts as the LNS for multiple LACs. The individual LACs do not have to have the administrative control or capability required to identify the most appropriate LNS on which to terminate their sessions. The LTS performs that function is centralized in the LTS.
- The LTS acts as the LAC for multiple LNSs. When a new remote LAC is added to an ISP's network, the ISP does not have to reconfigure its LNS routers to accommodate the new LAC, because they connect to the LAC on the LTS.

In a Layer 2 wholesale network, the wholesaler can use L2TP tunnel switching to create a flatter network configuration that is easier to manage. The wholesaler bundles Layer 2 sessions from a LAC that are destined for different ISPs—and therefore different LNSs—onto a single L2TP tunnel. This configuration enables a common L2TP control connection to be used for the LAC.


Figure 11 on page 155 shows an example of L2TP tunnel switching for incoming calls with the following sequence of events:

1. The subscriber opens a PPP session to the LAC.
2. The LAC creates the first L2TP tunnel to the LNS configured on the LTS and the first L2TP session to carry the encapsulated PPP packets.
3. During authentication of this first session, the LTS determines whether to retunnel the session to an LNS beyond the LTS, based on the presence or absence of a tunnel switch profile configured on the LTS.

The tunnel switch profile can be a default profile or it can be applied by the RADIUS server, a domain map configuration, or a tunnel group configuration.

4. If a tunnel switch profile is configured, the LTS creates a second tunnel (if it does not already exist) to the LNS beyond the LTS as specified in the profile and creates the second session in this tunnel.

Figure 11: L2TP Tunnel Switching for Incoming Calls



Application of Tunnel Switch Profiles

You can configure a tunnel switch profile to be applied in several ways:

- As a default profile applied globally to traffic received from all LACs
- With a domain map applied to a subscriber session
- With a tunnel group applied to a subscriber session
- In your RADIUS server configuration, returned in the Tunnel Switch-Profile VSA (26-91)

You can configure more than one of these methods of application. When multiple tunnel switch profiles are present, the following order of precedence establishes which profile the LTS uses; the order is from highest (RADIUS) to lowest (default profile):

RADIUS VSA 26-91 > domain map > tunnel group > global tunnel switch profile

The tunnel switch profile must also reference a tunnel profile. This tunnel profile specifies the characteristics of the second tunnel, to which the subscriber packets are switched.

Termination of Tunnel-Switched Sessions on the LTS

Tunnel switched sessions are terminated on the LTS when any of the following happens:

- Either the LAC or LNS interface on the LTS receives a Call-Disconnect-Notify (CDN) message ([Table 10 on page 156](#)).

Table 10: Cause of CDN Message

CDN Message Is Received On	When
LAC interface	Either of the following occurs: <ul style="list-style-type: none"> • The second session cannot be established. • The remote LNS terminates the second session.
LNS interface	Either of the following occurs: <ul style="list-style-type: none"> • The PPPoE client initiates a logout. • The originating LAC initiates termination of the tunnel

Both the first and second sessions are terminated because the LTS relays the CDN to the interface that did not receive the CDN. The disconnect cause is the same for both sessions.

- Either the LAC or LNS interface on the LTS receives a Stop-Control-Connection-Notification (StopCCN) message ([Table 11 on page 156](#)).

Table 11: Cause of StopCCN Message

StopCCN Message Is Received On	When
LAC interface	Either of the following occurs: <ul style="list-style-type: none"> • The second session cannot be established. • The remote LNS terminates the second tunnel.
LNS interface	The originating LAC initiates termination of the tunnel.

The LTS does not relay the StopCCN message, because a given tunnel can contain both switched and nonswitched sessions. Another reason in a wholesale scenario is that the tunnel ending on the LNS on the LTS can contain sessions from LACs from different providers. Instead, the LTS sends a CDN message to the interface that did not receive the StopCCN to terminate the tunnel-switched session. This CDN relays the error code carried in the StopCCN.

- An administrative **clear** command is issued on the LTS.

[Table 12 on page 156](#) lists the actions taken when an administrative **clear** command is issued on the LTS.

Table 12: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands

Command	LAC or LNS Action	LTS Action
clear services l2tp destination	Clear the destination and all associated tunnels and sessions.	For each switched session in a tunnel to the destination, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.

Table 12: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands (*continued*)

Command	LAC or LNS Action	LTS Action
<code>clear services l2tp destination all</code>	Clear all destinations and all associated tunnels and sessions.	None.
<code>clear services l2tp session</code>	Clear the session.	Clear the corresponding mapped switched session for this session by sending it a CDN message with the cause set to Administrative.
<code>clear services l2tp session all</code>	Clear all sessions.	None.
<code>clear services l2tp tunnel</code>	Clear the tunnel and all its sessions.	For each switched session in the tunnel, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.
<code>clear services l2tp tunnel all</code>	Clear all tunnels.	None.

Related Documentation

- [Configuring L2TP Tunnel Switching on page 161](#)
- [Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 157](#)
- [L2TP for Subscriber Access Overview on page 147](#)

Tunnel Switching Actions for L2TP AVPs at the Switching Boundary

When L2TP tunnel switching redirects packets to a different LNS, it performs one of the following default actions at the switching boundary for each AVP carried in the L2TP messages:

- **relay**—L2TP transparently forwards the AVP in the switched packet with no alteration.
- **regenerate**—L2TP ignores the received AVP that was negotiated by the first tunnel and session. It generates a new AVP for the second session based on the local policy at the LTS and sends this AVP in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.

[Table 13 on page 158](#) lists the default action for each AVP. Mandatory AVPs are always included in the L2TP messages from the LAC; optional AVPs might be included in the messages.

You can optionally override the default action taken at the switching boundary for the Bearer Type AVP (18), Calling Number AVP (22), or Cisco NAS Port Info AVP (100). You can configure any of these three AVPs to be dropped from the switched packets or regenerated, or you can restore the default relay action.



NOTE: L2TP AVPs that have their attribute values hidden are always regenerated at the switching boundary. The value is decoded and sent in clear text when the packet is forwarded to the remote LNS.

Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
Assigned Session Id (14)	Mandatory	CDN, ICRQ	Regenerate
Assigned Tunnel Id (9)	Mandatory	SCCRQ	Regenerate
Bearer Capabilities (4)	Optional	SCCRQ	Regenerate
Bearer Type (18)	Optional	ICRQ	Relay
Call Serial Number (15)	Mandatory	ICRQ	Relay
Called Number (21)	Optional	ICRQ	Relay
Calling Number (22)	Optional	ICRQ	Relay
Challenge (11)	Optional	SCCRQ	Regenerate
Challenge Response (13)	Optional	SCCCN	Regenerate
Cisco NAS Port	Optional	ICRQ	Relay
Failover Capability	Optional	SCCRQ	Regenerate
Firmware Revision (6)	Optional	SCCRQ	Regenerate
Framing Capabilities (3)	Mandatory	SCCRQ	Regenerate
Framing Type (19)	Mandatory	ICCN	Relay
Host Name (7)	Mandatory	SCCRQ	Regenerate
Initial Received LCP CONFREQ (26)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, the AVP is regenerated rather than relayed.

Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary (*continued*)

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
Last Received LCP CONFREQ (28)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, the AVP is regenerated rather than relayed.
Last Sent LCP CONFREQ (27)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, the AVP is regenerated rather than relayed.
Message Type (0)	Mandatory	All	Regenerate
Physical Channel Id (25)	Optional	ICRQ	Regenerate
Private Group Id (37)	Optional	ICCN	Relay
Protocol Version (2)	Mandatory	SCCRQ	Regenerate
Proxy Authen Challenge (31)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, authentication is also renegotiated and the AVP is regenerated rather than relayed.
Proxy Authen ID (32)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, authentication is also renegotiated and the AVP is regenerated rather than relayed.

Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary (*continued*)

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
Proxy Authen Name (30)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, authentication is also renegotiated and the AVP is regenerated rather than relayed.
Proxy Authen Response (33)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, authentication is also renegotiated and the AVP is regenerated rather than relayed.
Proxy Authen Type (29)	Optional	ICCN	Relay When LCP renegotiation is enabled with the lcp-negotiation statement in the client profile on the LNS, authentication is also renegotiated and the AVP is regenerated rather than relayed.
Receive Window Size (10)	Optional	SCCRQ	Regenerate
Rx Connect Speed (38)	Optional	ICCN	Relay
Sequencing Required (39)	Optional	ICCN	Regenerate
Sub-Address (23)	Optional	ICRQ	Relay
Tie Breaker (5)	Optional	SCCRQ	Regenerate
Tunnel Recovery	Optional	SCCRQ	Regenerate
Tx Connect Speed (24)	Mandatory	ICCN	Relay
Vendor Name (8)	Optional	SCCRQ	Regenerate

Related Documentation • [L2TP Tunnel Switching Overview on page 153](#)

- [Configuring L2TP Tunnel Switching on page 161](#)

Configuring L2TP Tunnel Switching

L2TP tunnel switching enables a router configured as an LTS to forward PPP packets carried on one L2TP session to a second L2TP session terminated on a different LNS. To configure L2TP tunnel switching, you must define a tunnel switch profile and then assign that profile.

You can configure tunnel switch profiles for all sessions globally, all sessions in a tunnel group, all sessions in a domain or in your RADIUS server configuration to be returned in the RADIUS Tunnel Switch-Profile VSA (26-91). The order of precedence for tunnel switch profiles from various sources is as follows:

- RADIUS VSA 26-91 > domain map > tunnel group > global tunnel switch profile

To define an L2TP tunnel switch profile:

1. Create the profile.

```
[edit access]
user@host# edit tunnel-switch-profile profile-name
```

2. (Optional) Override the default actions taken for certain L2TP AVPs at the switching boundary.

```
[edit access tunnel-switch-profile profile-name]
user@host# set avp bearer-type action
user@host# set avp calling-number action
user@host# set avp cisco-nas-port-info action
```

3. Specify the tunnel profile that defines the tunnel to which the subscriber traffic is switched.



NOTE: This step is not required for a tunnel switch profile specified in the Tunnel Switch-Profile VSA (26-91).

```
[edit access tunnel-switch-profile profile-name]
user@host# set tunnel-profile profile-name
```

4. (Optional) Apply the profile as a global default profile to switch packets from all incoming sessions from the LAC.

```
[edit services l2tp]
user@host1# set tunnel-switch-profile profile-name
```

5. (Optional) Apply the profile as part of a tunnel group to switch packets from all sessions in the tunnel group.

```
[edit services l2tp tunnel-group name]
```

```
user@host1# set tunnel-switch-profile profile-name
```



NOTE: The tunnel group is part of the LTS configuration that enables it to act as the LNS for the original sessions from the LAC.

A tunnel group with a tunnel switch profile must also contain a dynamic profile, because tunnel switching supports only dynamic subscribers.

6. (Optional) Apply the profile as part of a domain map to switch packets from all sessions that are associated with the domain.

```
[edit access domain-map domain-map-name]
user@host1# set tunnel-switch-profile profile-name
```



NOTE: A domain map cannot have both a tunnel switch profile and a tunnel profile. You must remove one if you add the other.

7. (Optional) Apply the profile by means of the Tunnel-Switch-Profile VSA [26–91] in the RADIUS Access-Accept message returned when the session from the LAC is authenticated. Refer to the documentation for your RADIUS server to determine how to configure this method.



NOTE: A tunnel switch profile specified by a RADIUS server in the Tunnel Switch-Profile VSA (26-91) takes precedence over the tunnel switch profile specified in the CLI configuration. If the Tunnel-Group VSA (26-64) is received in addition to the Tunnel Switch-Profile VSA (26-91), the Tunnel Switch-Profile VSA (26-91) takes precedence over the Tunnel-Group VSA (26-64), ensuring that the subscribers are tunnel switched rather than LAC tunneled.

For example, consider the following configuration, which creates three tunnel switch profiles, l2tp-tunnel-switch-profile, lts-profile-groupA, and lts-profile-example-com:

```
[edit access tunnel-switch-profile l2tp-tunnel-switch-profile]
user@host# set avp bearer-type regenerate
user@host# set avp calling-number regenerate
user@host# set avp cisco-nas-port-info drop
user@host# set tunnel-profile l2tp-tunnel-profile1
```

```
[edit access tunnel-switch-profile lts-profile-groupA]
user@host# set tunnel-profile l2tp-tunnel-profile2
[edit access tunnel-switch-profile lts-profile-example.com]
user@host# set tunnel-profile l2tp-tunnel-profile3
```

```
[edit services l2tp]
```

```
user@host1# set tunnel-switch-profile l2tp-tunnel-switch-profile
user@host1# set tunnel-group groupA tunnel-switch-profile lts-profile-groupA
```

```
[edit access domain]
user@host1# set map example.com tunnel-switch-profile lts-profile-example.com
```

The profile `l2tp-tunnel-switch-profile` is applied as the global default. When packets are switched according to this profile, the values for the Bearer Type AVP (18) and Calling Number AVP (22) in the L2TP packets are regenerated based on local policy at the L2TP tunnel switch and then sent with the packets. The Cisco NAS Port Info AVP (100) is simply dropped. Finally, `l2tp-tunnel-profile1` provides the configuration characteristics of the tunnel to which the traffic is switched.

Tunnel switch profile `lts-profile-groupA` is applied by means of a tunnel group, `groupA`; it specifies a different tunnel profile, `l2tp-tunnel-profile2` and it does not override any AVP actions. Tunnel switch profile `lts-profile-example.com` is applied by means of a domain map for the `example.com` domain; it specifies a different tunnel profile, `l2tp-tunnel-profile3` and it does not override any AVP actions.

- Related Documentation**
- [L2TP Tunnel Switching Overview on page 153](#)
 - [Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 157](#)
 - [Specifying a Tunnel Switch Profile in a Domain Map](#)

Setting the L2TP Receive Window Size

You can configure the L2TP receive window size for an L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

By default, the receive window size is set to four packets. If the receive window size is set to its default value, the router does not send the Receive Window Size AVP, AVP 10, in its first packet sent during tunnel negotiation to its peer.

To configure the receive window size:

```
[edit services l2tp tunnel]
user@host# set rx-window-size packets
```

- Related Documentation**
- [Configuring an L2TP LAC on page 173](#)
 - [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Setting the L2TP Tunnel Idle Timeout

You can configure the LAC or the LNS to specify how long a tunnel without any sessions remains active. The idle timer starts when the last session on the tunnel is terminated. When the timer expires the tunnel is disconnected. This idle timeout frees up resources otherwise consumed by inactive tunnels.

If you set the idle timeout value to zero, the tunnel is forced to remain active indefinitely after the last session is terminated until one of the following occurs:

- You issue the **clear services l2tp tunnel** command.
- The remote peer disconnects the tunnel.



BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no idle-timeout** statement at the **[edit services l2tp tunnel]** hierarchy level.

To set the tunnel idle timeout:

- Configure the timeout period.

```
[edit services l2tp tunnel]  
user@host# set idle-timeout seconds
```

**Related
Documentation**

- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Setting the L2TP Destruct Timeout

You can configure the LAC or the LNS to specify how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed. This destruct timeout aids debugging and other analysis by saving underlying memory structures after the destination, tunnel, or session is terminated. Any specific dynamic destination, tunnel, or session may not be maintained for this entire time period if the resources must be reclaimed early to allow new tunnels to be established.



BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no destruct-timeout** statement at the **[edit services l2tp]** hierarchy level.

To set the L2TP destruct timeout:

- Configure the timeout period.

```
[edit services l2tp]  
user@host# set destruct-timeout seconds
```

**Related
Documentation**

- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Configuring the L2TP Destination Lockout Timeout

When multiple sets of tunneling parameters are available, L2TP uses a selection process to choose the best tunnel for subscriber traffic. As part of this selection process, L2TP locks out destinations it cannot connect to when a subscriber tries to reach a domain. L2TP places the destination on the destination lockout list and excludes the destination from consideration for a configurable period called the *destination lockout timeout*.

By default, the destination lockout timeout is 300 seconds (5 minutes). You can configure a value from 60 through 3600 seconds (1 minute through 1 hour). When the lockout timeout expires, L2TP assumes that the destination is now available and includes the destination when performing the tunnel selection process. The destination lockout period is a global value and is not individually configurable for particular destinations, tunnels, or tunnel groups.



BEST PRACTICE: Configure the lockout timeout to be equal to or shorter than the destruct timeout. Otherwise, the destruct timeout expires before the lockout timeout. In this event, the locked-out destination is destroyed and can be subsequently returned to service before the lockout timeout expires, thus negating the effectiveness of the lockout timeout.

To configure the destination lockout timeout:

- Specify the period in seconds.

```
[edit services l2tp destination]
user@host# set lockout-timeout seconds
```

The **show services l2tp destination lockout** command displays the destination lockout list and for each destination indicates how much time remains before its timeout expires. The **show services l2tp destination detail** command indicates for each destination whether it is locked and waiting for the timeout to expire or not locked.

Related Documentation

- [LAC Tunnel Selection Overview on page 183](#)
- [Setting the L2TP Destruct Timeout on page 164](#)
- [Removing an L2TP Destination from the Destination Lockout List on page 165](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Removing an L2TP Destination from the Destination Lockout List

When a PPP subscriber tries to log in to a domain, L2TP selects a tunnel associated with a destination in that domain and attempts to access the destination. If the connection attempt fails, L2TP places the destination on the destination lockout list. Destinations on this list are excluded from being considered for subsequent connections for a configurable period called the *destination lockout timeout*.

You can issue the **request services l2tp destination unlock** command for a particular destination to remove it from the destination lockout list. The result is that this destination is immediately available for consideration when a subscriber logs in to the associated domain.

To remove a destination from the destination lockout list:

- Specify the name of the destination to be unlocked.

```
user@host> request services l2tp destination unlock destination-name
```

**Related
Documentation**

- [LAC Tunnel Selection Overview on page 183](#)
- [Configuring the L2TP Destination Lockout Timeout on page 165](#)

Configuring L2TP Drain

For administrative purposes, you can set the state of an L2TP destination or tunnel to drain. This prevents the creation of new sessions, tunnels, and destinations at L2TP LAC and LNS.

You can configure L2TP drain at the global level or for a specific destination or tunnel. If the feature is configured at global L2TP level, then no new destination, tunnel, or session can be created. If the feature is configured for a specific destination, no new tunnel or session can be created at that destination. Similarly, if the feature is configured for a specific tunnel, no new sessions can be assigned to that tunnel, but new destinations and tunnels can be created.

- To prevent creation of new sessions, destinations, and tunnels for L2TP:

```
[edit services]  
user@host# set l2tp drain
```

- To prevent creation of new tunnels and sessions at a particular destination:

```
[edit services]  
user@host# set l2tp destination address ip-address drain  
user@host# set l2tp destination address ip-address routing-instance  
                  routing-instance-name drain  
user@host# set l2tp destination name name drain
```

- To prevent creation of new sessions at a specific tunnel:

```
[edit services]  
user@host# set l2tp tunnel name name drain  
user@host# set l2tp tunnel name name address ip-address drain  
user@host# set l2tp tunnel name name address ip-address routing-instance  
                  routing-instance-name drain
```



NOTE: The tunnel *name* is the locally assigned name of the tunnel in the following format:

destination-name/tunnel-name or *tunnel-name*

When only the *tunnel-name* is provided, then you must include the address *ip-address* statement to identify the destination for the tunnel by.

When this feature is configured, the command output of **show services l2tp summary**, **show services l2tp destination**, and **show services l2tp tunnel** displays the state of the L2TP session, destination, and tunnel as **Drain**.

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Using the Same L2TP Tunnel for Injection and Duplication of IP Packets

You can configure the same L2TP tunnel that is used for subscriber secure policy mirroring to be used for duplication of packets. Packets duplicated are used to inject traffic towards the customer or towards the network. Injection or transmission of packets is supported for all subscriber access modes. A single L2TP tunnel is used for both transmission of packets and duplication of packets. A port or interface that is configured for duplication of packets on one side of an L2TP tunnel is connected to the other tunnel endpoint. The other endpoint of the tunnel can send IP packets using the L2TP tunnel to the port or interface configured for packet duplication, and the IP packets received at that interface can be either forwarded to the customer or sent as though it has been received from the customer.

The remote tunnel endpoint sends an IP tunnel packet that contains an Ethernet MAC address in the payload. If the destination MAC address of the payload packet contains the MAC address of the router, the Ethernet packet is sent in the outgoing direction towards the network, and it is processed and forwarded as though it is received on the customer port. If the source MAC address of the payload packet contains the MAC address of the router, the Ethernet packet is transmitted in the outgoing direction towards the customer port. If the tunnel does not contain the receive-cookie configured, packet injection does not happen. In such a case, any received tunnel packet is counted and dropped in the same manner in which packets that arrive with a wrong cookie are counted and dropped.

To configure the packet to be duplicated and sent towards the customer or the network (based on the MAC address in the Ethernet payload), include the **decapsulate l2tp output-interface *interface-name* cookie l2tpv3-cookie** statement at the **[edit firewall family *family-name* filter *filter-name* term *term-name* then]** hierarchy level. You can also configure a counter for the duplicated or decapsulated L2TP packets by including the **count counter-name** statement at the **[edit firewall family *family-name* filter *filter-name* term *term-name* then]** hierarchy level.

Configuring L2TP Control Messages for Subscribers

- [Retransmission of L2TP Control Messages on page 169](#)
- [Configuring Retransmission Attributes for L2TP Control Messages on page 171](#)

Retransmission of L2TP Control Messages

L2TP peers maintain a queue of control messages that must be sent to the peer device. After the local peer (LAC or LNS) sends a message, it waits for a response from the remote peer. If a response is not received, the local peer retransmits the message. This behavior allows the remote peer more time to respond to the message.

You can control the retransmission behavior in the following two ways:

- **Retransmission count**—You can configure how many times an unacknowledged message is retransmitted by the local peer. Increasing the count provides more opportunities for the remote peer to respond, but also increases the amount of control traffic. For tunnels that have been established, include the **retransmission-count-established** statement at the **[edit services l2tp tunnel]** hierarchy level. For tunnels that are not yet established, include the **retransmission-count-not-established** statement.
- **Retransmission interval**—You can configure how long the local peer waits for the first response to a control message. If a response is not received within the first timeout interval, then the retransmission timer doubles the interval between each successive retransmission up to a maximum of 16 seconds. Increasing the interval gives the remote peer more time to respond, but also spends more resources on a potentially unavailable peer. Include the **minimum-retransmission-interval** statement at the **[edit services l2tp tunnel]** hierarchy level.

The local peer continues retransmitting the control message until one of the following occurs:

- A response is received within the current waiting period.
- The maximum retransmission count is reached.

If the maximum count is reached and no response has been received, the tunnel and all its sessions are cleared.



NOTE: Reaching the maximum interval of 16 seconds does not halt retransmissions. The local peer continues to wait 16 seconds after each subsequent retransmission.

The following examples describe the retransmission behavior in different circumstances:

- Example 1—The retransmission count is three and the minimum retransmission interval is 1 second.
 1. The local peer sends a control message.
 2. The local peer waits 1 second, but receives no response.
 3. The local peer retransmits the control message. This is the first retransmission.
 4. The local peer waits 2 seconds, but receives a response before the interval expires.
 5. Retransmission stops because a response is received within the interval.
- Example 2—The retransmission count is two and the minimum retransmission interval is 8 seconds.
 1. The local peer sends a control message.
 2. The local peer waits 8 seconds, but receives no response.
 3. The local peer retransmits the control message. This is the first retransmission.
 4. The local peer waits 16 seconds, but receives no response.
 5. The local peer retransmits the control message. This is the second retransmission.
 6. The local peer again waits 16 seconds, because the interval cannot increase beyond 16, but receives no response.
 7. Retransmission stops because the maximum retransmission count of two was reached.
 8. The tunnel and all its sessions are cleared.

**Related
Documentation**

- [Configuring Retransmission Attributes for L2TP Control Messages on page 171](#)
- [L2TP for Subscriber Access Overview on page 147](#)

Configuring Retransmission Attributes for L2TP Control Messages

You can control the retransmission of unacknowledged L2TP control messages by configuring how many times the local peer retransmits the message and how long it waits for a response before retransmission.

L2TP peers maintain a queue of control messages that must be sent to the peer device. After the local peer (LAC or LNS) sends a message, it waits for a response from the remote peer. If a response is not received within the minimum retransmission interval, the local peer retransmits the message and waits for double the retransmission interval. Each time it retransmits the message, the peer doubles how long it waits, up to a maximum of 16 seconds.

If no response is received, the local peer continues to send the message until the number of retransmissions matches the retransmission count. In this case, retransmissions stop and the tunnel and all its sessions are cleared.



BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support these statements, we recommend that you explicitly unconfigure the feature by including the `no retransmission-count-established` statement and the `no retransmission-count-non-established` statement at the `[edit services l2tp tunnel]` hierarchy level.



BEST PRACTICE: During a unified in-service software upgrade (unified ISSU) on an MX Series router configured as the LAC, the LAC does not respond to control messages from the LNS. This can result in dropping LAC L2TP sessions. You can avoid this situation by ensuring that the maximum retransmission count on the LNS is set to 16 or higher.

To set the maximum retransmission count for established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-established count
```

To set the maximum retransmission count for non-established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-not-established count
```

To set the minimum interval between retransmissions:

- Configure the interval.

```
[edit services l2tp tunnel]
user@host# set minimum-retransmission-timeout seconds
```

For example, the following configuration specifies that established tunnels have a maximum retransmission count of three and a minimum retransmission interval of two seconds:

```
[edit services l2tp tunnel]
user@host# set retransmission-count-established 3
user@host# set minimum-retransmission-timeout 2
```

With this sample configuration, the following sequence applies to each control message sent by the LAC or LNS:

1. The local peer sends the control message and waits for a response from the remote peer.
2. If the response is not received within the minimum interval of 2 seconds, the local peer retransmits the message. This is the first retransmission.
3. If the response is not received within 4 seconds, the local peer retransmits the message. This is the second retransmission.
4. If the response is not received within 8 seconds, the local peer retransmits the message. This is the third and final retransmission, because the maximum count has been reached.
5. If the response is not received within 16 seconds, the tunnel and all its sessions are cleared.

**Related
Documentation**

- [Retransmission of L2TP Control Messages on page 169](#)
- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

CHAPTER 20

Configuring L2TP LAC Subscribers

- [Configuring an L2TP LAC on page 173](#)
- [Configuring How the LAC Responds to Address and Port Changes Requested by the LNS on page 175](#)
- [LAC Interoperation with Third-Party LNS Devices on page 177](#)
- [Globally Configuring the LAC to Interoperate with Cisco LNS Devices on page 178](#)
- [Configuring Username Modification for Subscriber Sessions on page 179](#)

Configuring an L2TP LAC

To configure an L2TP LAC:

1. Configure a tunnel profile to apply to subscribers.
[See “Configuring a Tunnel Profile for Subscriber Access” on page 207.](#)
2. (Optional) Configure the method used for selecting among multiple tunnels.
 - [See “Configuring the L2TP LAC Tunnel Selection Parameters” on page 209.](#)
 - [See “Configuring Weighted Load Balancing for LAC Tunnel Sessions” on page 211.](#)
 - [See “Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions” on page 211.](#)
 - [See “Configuring LAC Tunnel Selection Failover Within a Preference Level” on page 210.](#)
3. (Optional) Configure the LAC to not send Calling Number AVP 22 to the LNS.
[See “Preventing the LAC from Sending Calling Number AVP 22 to the LNS” on page 235.](#)
4. (Optional) Specify the method for setting the transmit and receive connect speeds.
[See “Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS” on page 228.](#)
5. (Optional) Configure whether the L2TP failover protocol is negotiated or the silent failover method is used for resynchronization.

See [“Configuring the L2TP Peer Resynchronization Method” on page 295.](#)

6. (Optional) Specify the format for the tunnel name.

See [“Setting the Format for the Tunnel Name” on page 206.](#)

7. (Optional) Specify when and how many times L2TP retransmits unacknowledged control messages.

See [“Configuring Retransmission Attributes for L2TP Control Messages” on page 171.](#)

8. (Optional) Specify how long a tunnel can remain idle before being torn down.

See [“Setting the L2TP Tunnel Idle Timeout” on page 163.](#)

9. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

See [“Setting the L2TP Receive Window Size” on page 163.](#)

10. (Optional) Specify how long the router retains information about terminated dynamic tunnels, sessions, and destinations.

See [“Setting the L2TP Destruct Timeout” on page 164.](#)

11. (Optional) Specify how the LAC handles IP address or UDP port change requests.

See [“Configuring How the LAC Responds to Address and Port Changes Requested by the LNS” on page 175.](#)

12. (Optional) Configure all tunnels on the LAC for interoperoperation with Cisco LNS devices.

See [“Globally Configuring the LAC to Interoperate with Cisco LNS Devices” on page 178.](#)

13. (Optional) Specify that the LAC sends information to the LNS about subscriber access lines.

See [“Configuring the Reporting and Processing of Subscriber Access Line Information” on page 230.](#)

14. (Optional) Configure the LAC to create the IPv6 address family (inet6) when establishing a tunnel for subscribers, enabling the application of IPv6 firewall filters.

See [“Enabling the LAC for IPv6 Services” on page 212.](#)

15. (Optional) Prevent the creation of new sessions, destinations, or tunnels for L2TP.

See [“Configuring L2TP Drain” on page 166.](#)

16. (Optional) Enable SNMP statistics counters.

See “Enabling Tunnel and Global Counters for SNMP Statistics Collection” on page 303.

17. (Optional) Configure trace options for troubleshooting the configuration.

See “Tracing L2TP Operations for Subscriber Access” on page 371.

Related Documentation

- [L2TP for Subscriber Access Overview on page 147](#)

Configuring How the LAC Responds to Address and Port Changes Requested by the LNS

An LNS can use the SCCRP message that it sends the LAC when a tunnel is being established to request a change in the destination IP address or UDP port that the LAC uses to communicate with the LNS. By default, the LAC accepts the request and makes the change. You can use the **tx-address-change** statement to configure one of the following methods for the LAC to handle these change requests for all tunnels:

- **accept**—The LAC accepts the change from the LNS. It sends all subsequent packets to and receives packets from the new IP address or UDP port.
- **ignore**—The LAC continues to send packets to the original address or port, but accepts packets from the new address or port.
- **reject**—The LAC sends a StopCCN message to the original address or port and then terminates the connection to that LNS.

The LAC accepts a change in address or port only once, when the tunnel is being established. Tunnels that are already established are not affected. The LAC drops any L2TP control packets containing change requests received at any other time, or in any packet other than an SCCRP message.



NOTE: This statement does not support IPv6 addresses.

To configure how the LAC handles change requests for the IP address, the UDP port, or both:

- (Optional) Configure the LAC to accept all change requests. This is the default behavior.

```
[edit services l2tp tunnel]
user@host# set tx-address-change accept
```

- (Optional) Configure the LAC to ignore all change requests.

```
[edit services l2tp tunnel]
user@host# set tx-address-change ignore
```

- (Optional) Configure the LAC to ignore change requests only for the IP address.

```
[edit services l2tp tunnel]
user@host# set tx-address-change ignore-ip-address
```

- (Optional) Configure the LAC to ignore change requests only for the UDP port.

```
[edit services l2tp tunnel]
user@host# set tx-address-change ignore-udp-port
```

- (Optional) Configure the LAC to reject all change requests.

```
[edit services l2tp tunnel]
user@host# set tx-address-change reject
```

- (Optional) Configure the LAC to reject change requests only for the IP address.

```
[edit services l2tp tunnel]
user@host# set tx-address-change reject-ip-address
```

- (Optional) Configure the LAC to reject change requests only for the UDP port.

```
[edit services l2tp tunnel]
user@host# set tx-address-change reject-udp-port
```

For example, the following configuration causes the LAC to ignore requests to change the UDP port, but to reject requests to change the IP address:

```
[edit services l2tp tunnel]
user@host# set tx-address-change ignore-udp-port
user@host# set tx-address-change reject-ip-address
```



NOTE: Conflicting configurations are not allowed and fail the configuration commit check. You cannot For example, the following configuration fails, because it specifies that UDP port changes are ignored, but that *all* changes are rejected:

```
[edit services l2tp tunnel]
user@host# set tx-address-change ignore-udp-port
user@host# set tx-address-change reject
```

Use the **show services l2tp summary** command to display the current behavior of the LAC:

```
show services l2tp summary
Failover within a preference level is Disabled
Weighted load balancing is Disabled
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 Disabled
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Ignore
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
Destinations: 1, Tunnels: 0, Sessions: 0
```

Depending on the configuration, this command displays one of the following outputs:

```
Tunnel Tx Address Change is Accept
Tunnel Tx Address Change is Ignore
Tunnel Tx Address Change is Reject
Tunnel Tx Address Change is Ignore IP Address & Accept UDP Port
Tunnel Tx Address Change is Ignore IP Address & Reject UDP Port
Tunnel Tx Address Change is Accept IP Address & Ignore UDP Port
Tunnel Tx Address Change is Accept IP Address & Reject UDP Port
Tunnel Tx Address Change is Reject IP Address & Accept UDP Port
Tunnel Tx Address Change is Reject IP Address & Ignore UDP Port
```

Related Documentation

- [Configuring an L2TP LAC on page 173](#)

LAC Interoperation with Third-Party LNS Devices

In some network environments, the LAC may need to interoperate with an LNS configured on a device from another vendor that does not run Junos OS. Interoperation with Cisco Systems devices requires the LAC to communicate a NAS port type, but the LAC does not provide this information by default.

You can enable interoperation with Cisco Systems devices by configuring the NAS port method as **cisco-avp**, which causes the LAC to include the Cisco Systems NAS Port Info AVP (100) when it sends an incoming call request (ICRQ) to the LNS. The AVP includes information that identifies the NAS port and indicates whether the port type is ATM or Ethernet.

You can configure the NAS port method globally for all tunnels on the LAC or in a tunnel profile for only the tunnels instantiated by the profile.

You can also include the Tunnel-Nas-Port-Method VSA [26–30] in your RADIUS server configuration with the value set to 1 to indicate Cisco Systems CLID. In this case, RADIUS can override the global value by modifying or creating a tunnel profile. The RADIUS configuration has precedence over the tunnel profile configuration, which in turn has precedence over the global LAC configuration.

If the LNS receiving the AVP is an MX Series router instead of a Cisco Systems device, the LNS simply ignores the AVP, unless the LNS is configured for L2TP tunnel switching. In that case, the LNS preserves the value of the AVP and passes it along when it switches tunnels for the LAC.

Related Documentation

- [Globally Configuring the LAC to Interoperate with Cisco LNS Devices on page 178](#)
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)
- [Configuring an L2TP LAC on page 173](#)
- *Juniper Networks VSAs Supported by the AAA Service Framework*

- [L2TP for Subscriber Access Overview on page 147](#)

Globally Configuring the LAC to Interoperate with Cisco LNS Devices

Cisco LNS devices require from the LAC both the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM. By default, the LAC does not include this information. You can globally configure the LAC to provide this information by including the NAS Port Info AVP (100) in the ICRQ that it sends to the LNS. This configuration enables the LAC to interoperate with a Cisco LNS.

To globally configure the LAC to include the NAS Port Info AVP:

- Specify the NAS port method.

```
[edit services l2tp tunnel]
user@host# set nas-port-method cisco-avp
```



NOTE: This global configuration for the LAC can be overridden by the configuration in a tunnel profile or RADIUS.

Use the **show services l2tp tunnel extensive** command to display the current behavior of the LAC:

```
show services l2tp tunnel extensive
Tunnel local ID: 51872, Tunnel remote ID: 8660
Remote IP: 192.0.2.20:1701
Sessions: 5, State: Established
Tunnel Name: 1/tunnel-test-2
Local IP: 203.0.113.2:1701
Local name: testlac, Remote name: ce-lns
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: default
Max sessions: 128100, Window size: 4, Hello interval: 60
Create time: Thu Jul 25 12:55:41 2013, Up time: 11:18:14
Idle time: 00:00:00
Statistics since: Thu Jul 25 12:55:41 2013
```

	Packets	Bytes
Control Tx	702	15.5k
Control Rx	690	8.5k
Data Tx	153.3k	6.6M
Data Rx	126.3k	5.9M
Errors Tx	0	
Errors Rx	0	

Related Documentation

- [LAC Interoperation with Third-Party LNS Devices on page 177](#)
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)
- [Configuring an L2TP LAC on page 173](#)

Configuring Username Modification for Subscriber Sessions

You can use subscriber session options to set parameters that modify a subscriber's username at login based on the subscriber's access profile. This modification is also called username *stripping*, because some of the characters in the username are stripped away and discarded. The remainder of the string becomes the new, modified username. The modified username is used by an external AAA server for session authentication and accounting. This capability can be useful, for example, in Layer 2 wholesale implementations, where the network service providers employ username modification to direct subscribers to the appropriate retail enterprise network.

The modification parameters are applied according to a subscriber access profile that also determines the subscriber and session context; that is, the logical system:routing instance (LS:RI) used by the subscriber. Only the default (master) logical system is supported. Because the wholesaler differentiates between multiple retailers by placing each in a different LS:RI, the usernames are appropriately modified for each retailer.

You can select up to eight characters as delimiters to mark the boundary between the discarded and retained portions of the original username; there is no default delimiter. The portion of the name to the right of the selected delimiter is discarded along with the delimiter. By configuring multiple delimiters, a given username structure can result in different modified usernames. You can configure the direction in which the original name is parsed to determine which delimiter marks the boundary. By default, the parse direction is from left to right.

To configure username modification:

1. Define a profile consisting of a set of AAA options for authorizing and configuring a subscriber or set of subscribers with a subscriber access profile.

- a. Specify the name of the subscriber access profile that includes the username stripping configuration.

```
[edit access aaa-options aaa-options-name]
user@host# access-profile profile-name
```

- b. (Optional) Specify the logical-system:routing-instance (LS:RI) that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting. For example, this may correspond to the LS:RI for a retail ISP that provides services to the subscriber.

```
[edit access aaa-options aaa-options-name]
user@host# aaa-context aaa-context-name
```

- c. (Optional) Specify the logical-system:routing-instance (LS:RI) in which the subscriber interface is placed. For example, this may correspond to the LAC-facing interface on the LNS that is accessed by all requests from a subscriber residence.

```
[edit access aaa-options aaa-options-name]
```

```
user@host# subscriber-context subscriber-context-name
```

2. Configure the session options in the access profile that specify how usernames are stripped.
 - a. Specify one or more delimiters to mark the boundary between the discarded and retained portions of the original username.

```
[edit access profile profile-name session-options strip-user-name]  
user@host# set delimiter [ delimiter ]
```

- b. (Optional) Specify the direction in which the original username is examined to find a delimiter. The default direction is left-to-right.

```
[edit access profile profile-name session-options strip-user-name]  
user@host# set parse-direction (left-to-right | right-to-left)
```

3. (Optional) Specify that the AAA options are on a per-interface basis when dynamic subscribers are authenticated.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit" ppp-options]  
user@host# set aaa-options aaa-options-name
```

4. (Optional) Specify that the AAA options are part of the PPP options in a group profile that applies to tunneled PPP subscribers at the LNS.

```
[edit access group-profile profile-name ppp]  
user@host# set ppp-options aaa-options aaa-options-name
```



In the following example, the AAA options profile, `aaa1`, specifies a subscriber access profile, `entA`, for subscribers in the default logical system and routing instance 1. The access profile, `entA`, specifies that usernames are examined from left to right until the delimiter, `@`, is found. The AAA options profile is applied to tunneled PPP subscribers that belong to the group profile, `FD1`.

```
[edit access aaa-options aaa1]  
user@host# access-profile entA  
user@host# aaa-context master:1
```

```
[edit access profile entA session-options strip-user-name]  
user@host# set delimiter @  
user@host# set parse-direction left-to-right
```

```
[edit access group-profile FD1 ppp]  
user@host# set ppp-options aaa-options aaa1
```

Given that configuration, suppose a subscriber attempts to log in with the username, `user1@example.com`. When this name is examined, the delimiter and the string `example.com` are discarded, leaving a modified username of `user1`. Note that the result is the same if the parse direction is set to examine the name from right to left, because only one delimiter is defined and only one is present in the original username.

parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1



8043376

Now suppose the subscriber logs in with the username, `user1@test@example.com`. For a username like this, the parsing direction makes a difference in the modified username. The configuration determines that the first instance of the delimiter `@` is found first, because the name is parsed from left to right. This delimiter and the string `test@example.com` are discarded, leaving `user1` as the modified username.

What happens when the configuration sets a different parsing direction?

```
[edit access profile entA session-options strip-user-name]
user@host# set delimiter @
user@host# set parse-direction right-to-left
```

In this case, for the username `user1@test@example.com`, the second instance of the delimiter is identified and it is discarded with the string `@example.com`. The modified username is `user1@test`.



parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1@test

8043377

You can achieve the same results of different modified usernames based on parse direction by configuring more than one delimiter as in the following configuration, where you specify two delimiters, `@` and `/`.

```
[edit access profile entA session-options strip-user-name]
user@host# set delimiter [ @ / ]
user@host# set parse-direction left-to-right
```

For the username `user1@bldg1/example.com`, parsing left to right identifies the `@` delimiter first and the modified username is `user1`. Parsing right to left instead, identifies the `/` delimiter first and strips it away with the string `example.com`, leaving a modified username of `user1@bldg1`.

parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1@bldg1

8043378

Related Documentation

- *Understanding Session Options for Subscriber Access*

- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)
- *Configuring Subscriber Session Timeout Options*

CHAPTER 21

Configuring L2TP LAC Tunneling for Subscribers

- [LAC Tunnel Selection Overview on page 183](#)
- [L2TP Session Limits Overview on page 199](#)
- [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204](#)
- [Setting the Format for the Tunnel Name on page 206](#)
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)
- [Configuring the L2TP LAC Tunnel Selection Parameters on page 209](#)
- [Configuring LAC Tunnel Selection Failover Within a Preference Level on page 210](#)
- [Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 211](#)
- [Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions on page 211](#)
- [Enabling the LAC for IPv6 Services on page 212](#)

LAC Tunnel Selection Overview

When a user logs in to a domain, the PPP client contacts the LAC to establish a connection. The LAC has to find a destination in the domain and a tunnel that can reach it. The association between destinations, tunnels, and domains is provided by a tunnel profile either in a domain map in the subscriber's access profile or in the Tunnel-Group attribute (VSA 26-64) received from a RADIUS server. The RADIUS attribute takes precedence over a profile specified in a domain map. The tunnel profile includes a list of tunnels; each tunnel is associated with a destination IP address and with a tunnel preference level.

L2TP enables you to specify:

- Up to 31 destinations for a domain.
- Up to eight levels of tunnel preference. The preference level determines the order in which the LAC attempts to use an existing tunnel (or establish a new one) to a destination in the user's requested domain.



NOTE: Zero (0) is the highest level of preference; this is the most-preferred level.

If two tunnels both reach valid destinations within a domain, the LAC first selects the tunnel with the highest preference level. For example, when Tunnel A has a preference level of 1 and Tunnel B has a preference level of 4, the LAC attempts to use Tunnel A first.

- Up to 31 destinations for a single preference level.

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from the set of tunnels associated with either the PPP user or the PPP user's domain by a tunnel profile.

Tunnel selection is affected by the following configurations:

- Failover between preference levels—By default, when a tunnel to a valid destination is not selected within a preference level, the selection process fails over to the next level; that is, the LAC drops down to the next lower level to continue the search for a suitable tunnel. See [“Selection When Failover Between Preference Levels Is Configured” on page 186](#) for more information.
- Failover within a preference level—In this case, the LAC does not limit its attempts to establish a session to only a single tunnel at a preference level. If the attempt fails through the selected tunnel, the selection process fails over within that same level by selecting another suitable tunnel to a valid destination. The LAC continues its connection attempts within the level until no more tunnels to a valid destination are available at that level. Then the LAC drops down to the next lower level to continue the search. See [“Selection When Failover Within a Preference Level Is Configured” on page 191](#) for more information.
- Maximum sessions per tunnel—When the maximum number of sessions allowed per tunnel is configured, the LAC takes that setting into account during the tunnel selection process. The maximum number of sessions per tunnel can be configured by means of the RADIUS Tunnel-Max-Sessions VSA [26-33] or by including the **max-sessions** statement in a tunnel profile.

When a randomly selected tunnel has a current session count equal to its maximum session count, the LAC does not attempt to connect to a destination with that tunnel. Instead, it selects an alternate tunnel from the set of tunnels at that preference level that have valid destinations in the domain. If no such tunnels exist at the current preference level, the LAC drops to the next preference level to make the selection. This process is consistent, regardless of which failover scheme is currently running on the LAC.

When the maximum number of sessions is not configured for a tunnel, then that tunnel has no upper limit on the number of sessions it can support. By default, the maximum sessions value is 0 (zero), which allows unlimited sessions in the tunnel.

- Weighted load balancing—This balancing method uses a probability-based evaluation of tunnel weight to distribute sessions across tunnels. The LAC still selects tunnels randomly within a preference level, but on average the sessions are distributed across tunnels in relationship to the weight of the tunnels. The weight of a tunnel is determined by the tunnel's maximum session limit and the maximum session limits of the other

tunnels at the same preference level. See [“Weighted Load Balancing” on page 193](#) for more information.

- Destination-equal load balancing—This session-balancing method evaluates tunnels according to the number of sessions to the destination and the number of sessions carried by the tunnel in order to spread the session load equally among all tunnels. The tunnel with a destination that has the lowest session count is determined to have the lightest load. This process operates on tunnels at the highest available preference level. See [“Destination-Equal Load Balancing” on page 194](#) for more information.

Take the following information into consideration to understand the tunnel and destination selection process and failover:

- More than one tunnel may be able to reach a destination, and those tunnels can have the same preference level or different preference levels.
- The tunnel selected to establish the subscriber session may itself already be established, meaning that it has currently active sessions. Alternatively, the LAC might have to establish a new tunnel to the destination if no tunnel capable of reaching the destination is already established.
- A *valid* destination meets the following criteria:
 - It is reachable by a tunnel that has not met its maximum session limit.
 - It has not yet been contacted for the current subscriber login request.
 - It can be either locked or unlocked.
- A *locked* destination is one for which the destination lockout timer is running. Locked destinations are placed on a lockout list until the timer expires or is cleared (reset to zero). Destinations on the list cannot be contacted to establish a session.
- An *unlocked* destination is one for which the destination lockout timer is zero.
- When the LAC discovers valid destinations that are locked, it places them on the DestinationsLockedNotContacted list, which is different than the lockout list that includes all locked-out destinations. The DestinationsLockedNotContacted list includes only locked destinations that the LAC has not yet attempted to contact for the current, in-progress subscriber login. The DestinationsLockedNotContacted list does not include destinations that the LAC locks out after it has attempted and failed to establish a connection.
- You can use the **clear services l2tp destination lockout** command to manually clear all locked destinations or only locked destinations that match the specified local or remote gateway address. You might use the command if, for example, you want to clear a specific destination so that it gets priority within a preference level.
- The failover behavior that is part of the tunnel selection process applies only when the destination is unreachable for one of the following reasons:
 - The LNS fails to return an SCCRP message in response to the SCCRQ message from the LAC after the maximum number of retransmission attempts.

- The tunnel is established, but the LNS does not return an ICRP message in response to the ICRQ from the LAC after the maximum number of retransmission attempts.
- This failover behavior does not apply in the following circumstances:
 - The client terminates the connection.
 - The tunnel is established, but the LNS sends a CDN message while the LAC is attempting to establish the session with the LNS, resulting in the failure of the subscriber login attempt.

Selection When Failover Between Preference Levels Is Configured

When a user tries to log in to a domain in a default configuration—that is, when failover within a preference level and load balancing are not configured—the LAC searches for valid destinations to the requested domain, starting at the highest tunnel preference level. If no valid destination is found, or the attempt to connect to a destination fails, the LAC drops down to the next lower level to continue searching. The search process is the same for all levels except for the lowest:

1. The search begins by identifying tunnels with valid destinations at the preference level from among all the tunnels specified in the domain's tunnel profile.
2. All locked, valid destinations are placed on the DestinationsLockedNotContacted list. No attempt is made to contact any of these destinations.
3. From among the unlocked, valid destinations, the LAC selects one at random and attempts to connect through the associated tunnel; if the tunnel has no current sessions, then the LAC must establish the tunnel.



NOTE: Random selection is the default behavior. The behavior is different when weighted load balancing or destination-equal load balancing is configured. See [“Selection When Distributing the Session Load Across Multiple LNSs” on page 193](#) for information about load balancing.


- If the attempt is successful, the LAC reports the successful login to the PPP client. The LAC also clears all destinations on the DestinationsLockedNotContacted list.
 - If the LAC receives no response, it retries the attempt up to the maximum retry number. If the LAC exhausts the retries without receiving a reply, the attempt is considered unsuccessful and the LAC marks the destination as unreachable by locking out the destination. It places the destination on the lockout list and starts the destination lockout timer.
4. What the LAC does next depends on the current preference level.
 - If it is not the lowest preference level, then the LAC drops to the next lower preference level and continues the search process.
 - If it is the lowest preference level and the DestinationsLockedNotContacted list is not empty, then the LAC unlocks all destinations in the

DestinationsLockedNotContacted list and jumps back up to the highest preference level and restarts the search process.

- If it is the lowest preference level and the DestinationsLockedNotContacted list is empty—meaning that all valid destinations have been attempted—then the LAC reports a failed login to the PPP client.
5. When the valid destinations at one level are all locked, what the LAC does next depends on the current preference level.
 - If it is not the lowest preference level, then the LAC drops to the next lower preference level and continues the search process.
 - If it is the lowest preference level, the LAC selects the locked, valid destination with the shortest remaining lockout time. It clears the lockout timer and attempts to connect to the destination and establish a session.
 - If the attempt is successful, the LAC reports the successful login to the PPP client.
 - If the attempt fails and the DestinationsLockedNotContacted list is empty—meaning that all valid destinations have been attempted—then the LAC reports a failed login to the PPP client.
 - If the attempt fails and the DestinationsLockedNotContacted list is not empty, then the LAC unlocks all destinations in the DestinationsLockedNotContacted list, jumps back up to the highest preference level, and restarts the search process.
 6. When no valid destinations are present, what the LAC does next depends on the current preference level.
 - If it is not the lowest preference level, then the LAC drops to the next lower preference level and continues the search process.
 - If it is the lowest preference level and the DestinationsLockedNotContacted list is empty—meaning that all valid destinations have been attempted—then the LAC reports a failed login to the PPP client.
 - If it is the lowest preference level and the DestinationsLockedNotContacted list is not empty, then the LAC unlocks all destinations in the DestinationsLockedNotContacted list, jumps back up to the highest preference level, and restarts the process.
 7. The search and failover process cycles through the levels until either a session is established or all valid destinations have been attempted—no destinations remain on the DestinationsLockedNotContacted list—and the login fails.

Figure 12 on page 188 illustrates the possible conditions and decision points that determine the selection of a destination and corresponding tunnel for the default case, where failover occurs between tunnel preference levels.

Figure 12: Destination and Tunnel Selection Process with Failover Between Preference Levels



For example, suppose that the tunnel profile includes the following tunnels, each with a valid destination:

- Preference 0, Tunnel 1, 192.168.10.10
- Preference 1, Tunnel 2, 192.168.22.22

- Preference 1, Tunnel 3, 192.168.33.33
- Preference 2, Tunnel 4, 192.168.44.44

Failover within preference and load balancing are not configured.

When a PPP user tries to connect to the domain, the LAC acts as follows:

1. At the highest preference level, 0, the LAC selects Tunnel 1 because it is the only tunnel in the level with a valid destination. The LAC attempts to reach 192.168.10.10.
2. This connection attempt fails, so the LAC locks out 192.168.10.10. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
3. The LAC drops (fails over) to the next level, preference level 1, to reach a destination for the domain. The LAC randomly selects between 192.168.22.22 through Tunnel 2 and 192.168.33.33 through Tunnel 3. It selects 192.168.22.22 and attempts to connect through Tunnel 2.
4. The connection attempt to 192.168.22.22 fails, so the LAC locks out 192.168.22.22. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.



NOTE: Even though Tunnel 3 has an unlocked, valid destination, the LAC cannot now select that tunnel to reach 192.168.33.33, because the LAC can make only one attempt to reach a valid destination each time it searches in a level when the failover method is between preference levels.

5. The LAC drops to the final (lowest) level in this example, preference level 2. The LAC selects Tunnel 4 because it is the only tunnel in the level with a valid destination. The LAC attempts to reach 192.168.44.44.
6. The connection attempt to 192.168.44.44 also fails, so the LAC locks out 192.168.44.44. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
7. Because this is the lowest level, and the DestinationsLockedNotContacted list is empty, the LAC rejects the login request from the PPP client.

Destinations 192.168.10.10, 192.168.22.22, and 192.168.44.44 were locked out, but not added to the DestinationsLockedNotContacted list because the LAC locked them out after attempting to connect. Destination 192.168.33.33 was not contacted, but not added to the DestinationsLockedNotContacted list because it is not locked out.

8. The client tries to log in again and the LAC repeats the tunnel selection process, starting over at preference level 0 to check for an unlocked, valid destination, and cycling through the levels as needed.
9. At preference level 0, 192.168.10.10 is the only valid destination and is still locked out, so the LAC cannot attempt to connect to the destination. The LAC adds 192.168.10.10 to the DestinationsLockedNotContacted list and then drops to preference level 1.



NOTE: Remember that the destination lockout timer applies globally, so it persists across multiple subscriber logins. The DestinationsLockedNotContacted list applies only to a given subscriber login and does not persist. Even though the LAC contacted 192.168.10.10 for this subscriber, it was during a previous login attempt. In this login attempt, it cannot contact the destination because of the lockout, and consequently places the destination on the DestinationsLockedNotContacted list.

10. At preference level 1, 192.168.22.22 is still locked out, so the LAC adds 192.168.22.22 to the DestinationsLockedNotContacted list. 192.168.33.33 is still available. The LAC attempts to connect to 192.168.33.33 through Tunnel 3.
11. This connection attempt fails, so the LAC locks out 192.168.33.33. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires. The LAC drops to preference level 2.
12. 192.168.44.44 is still locked out, so the LAC adds 192.168.44.44 to the DestinationsLockedNotContacted list.
13. This is the lowest preference level, but this time the DestinationsLockedNotContacted list is not empty; it contains 192.168.10.10, 192.168.22.22, and 192.168.44.44. The LAC unlocks all destinations on the DestinationsLockedNotContacted list and then jumps back to the highest preference level.
14. At preference level 0, the LAC attempts to connect to 192.168.10.10 because it was unlocked. The LAC establishes the session and reports the successful login to the PPP client.

Although the LAC does not attempt to contact a destination that is locked out, there is a special case when the LAC has reached the lowest preference level. The level must have more than one valid destination and all of them must be locked out. For example, suppose that the tunnel profile includes the following tunnels, each with a valid destination:

- Preference 0, Tunnel 1, 192.168.10.10
- Preference 1, Tunnel 2, 192.168.22.22. The destination is locked out with the lockout timer currently at 245 seconds.
- Preference 1, Tunnel 3, 192.168.33.33. The destination is locked out with the lockout timer currently at 180 seconds.

Failover within preference and load balancing are not configured.

When a PPP user tries to connect to the domain, the LAC acts as follows:

1. At the highest preference level, 0, the LAC selects Tunnel 1 because it is the only tunnel in the level with a valid destination. The LAC attempts to reach 192.168.10.10.
2. This connection attempt fails, so the LAC locks out 192.168.10.10. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
3. The LAC drops to the next level, preference level 1, to reach a destination for the domain. Both valid destinations at this level, 192.168.22.22 and 192.168.33.33, are locked out.
4. The LAC adds both destinations to the DestinationsLockedNotContacted list.
5. Because this is the lowest preference level, the LAC determines which destination has a shorter remaining lockout time. It selects 192.168.33.33 because it has a shorter remaining lockout time (180 seconds) than 192.168.22.22 (245 seconds). The LAC unlocks 192.168.33.33 and attempts to connect through Tunnel 3. As a consequence, the LAC also removes 192.168.33.33 from the DestinationsLockedNotContacted list.
6. The connection attempt is successful and a session is established to 192.168.33.33. The LAC reports a successful login to the PPP client.

Selection When Failover Within a Preference Level Is Configured

When you configure failover *within* a preference level, the destination and tunnel selection process is the same as for the default configuration, with one exception: the LAC is not limited to only one connection attempt at a preference level. When the LAC tries to connect to an unlocked, valid destination and is unsuccessful, it locks out that destination but does not immediately drop down to the next lower level. Instead, if another unlocked, valid destination is available at the same preference level, the LAC attempts to connect to that destination. If the LAC does not connect, then it continues to try to reach a destination within that preference level until no more unlocked, valid destinations remain to be attempted. At that point the LAC drops down to search at the next lower preference level. At each level, the LAC searches for and attempts to connect to a valid destination until no unlocked, valid destinations are available.

For example, suppose that the tunnel profile specifies the following tunnels and destinations. Load balancing is not configured. All destinations are valid; all are unlocked except 192.168.3.3. The preference levels for the tunnels are assigned as follows:

- Preference 0, Tunnel 1, 192.168.1.1, unlocked
- Preference 0, Tunnel 2, 192.168.2.2, unlocked
- Preference 0, Tunnel 3, 192.168.3.3, lockout timer 100 seconds
- Preference 1, Tunnel 4, 192.168.4.4, unlocked
- Preference 1, Tunnel 5, 192.168.5.5, unlocked

In this example, when a PPP user tries to connect to the domain, the LAC acts as follows:

1. The LAC randomly selects between the two unlocked, valid destinations at preference level 0, 192.168.1.1 through Tunnel 1 and 192.168.2.2 through Tunnel 2. It chooses 192.168.2.2 and attempts to connect through Tunnel 2.
2. The connection attempt to 192.168.2.2 fails, so the LAC locks out 192.168.2.2. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
3. The LAC then attempts to connect to 192.168.1.1 through Tunnel 1 at preference level 0.
4. The connection attempt to 192.168.1.1 fails, so the LAC locks out 192.168.1.1. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
5. 192.168.3.3 through Tunnel 3 is the only remaining valid destination at preference level 0, but it is locked. The LAC adds 192.168.3.3 to the DestinationsLockedNotContacted list. The LAC did not add 192.168.1.1 and 192.168.2.2 to the DestinationsLockedNotContacted list, because it locked them out after attempting to contact them.
6. Because level 0 has no more unlocked, valid destinations, the LAC drops to the next level, preference level 1, to reach a destination for the domain.
7. At preference level 1, the LAC randomly selects 192.168.4.4 and attempts to connect through Tunnel 4.
8. The connection attempt to 192.168.4.4 fails, so the LAC locks out 192.168.4.4. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
9. The LAC then attempts to connect to 192.168.5.5 through Tunnel 5 at preference level 1.
10. The connection attempt to 192.168.5.5 fails, so the LAC locks out 192.168.5.5. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires. Level 1 has no more unlocked, valid destinations. Because the DestinationsLockedNotContacted list is not empty, the LAC unlocks all the destinations on the list—in this case, 192.168.3.3—and jumps back up to the highest preference level, 0.
11. 192.168.3.3 is now the only unlocked destination at preference level 0, so the LAC attempts to connect to it through Tunnel 3.
12. The connection attempt to 192.168.3.3 fails, so the LAC locks out 192.168.3.3. It is not considered again during this login attempt, and cannot be considered for any login attempt until the destination lockout timer expires.
13. Because level 0 has no more unlocked, valid destinations, the LAC drops to the next level, preference level 1.
14. Preference level 1 has no unlocked, valid destinations. The DestinationsLockedNotContacted is empty because the LAC has contacted all valid

destinations at both preference levels. The LAC rejects the login request from the PPP client.

Selection When Distributing the Session Load Across Multiple LNSs

Multiple tunnel profiles can be configured on the LAC; some tunnels may share destinations. When the LAC tunnels the session for a PPP subscriber to the LNS, a tunnel has to be selected for the subscriber session. The tunnel selection process chooses a tunnel with the highest preference that has a reachable destination. By default, the LAC selects a tunnel at random from among multiple tunnels that meet the same criteria. Alternatively, you can configure load balancing to enable different selection choices. Both load-balancing methods affect which tunnels and destinations the LAC selects, but the selection and failover process otherwise remains the same.



NOTE: Weighted load balancing and destination-equal load balancing are mutually exclusive. You can enable only one or the other.

Weighted Load Balancing

Weighted load balancing evaluates tunnels according to their weight. The weight of a tunnel is determined by the tunnel's maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the highest maximum session limit has the highest weight in that preference level. The tunnel with the next-highest maximum session limit has the next-highest weight, and so on. The tunnel with the lowest maximum session limit has the lowest weight.



NOTE: Tunnel selection and session distribution are probability based; the load is not strictly distributed according to weight.

When you configure weighted load balancing, the LAC still selects tunnels randomly within a preference level, but on average the sessions are distributed across tunnels in relationship to the weight of the tunnels.

With weighted load balancing, the LAC generates a random number within a range equal to the aggregate total of all session limits for all tunnels in the preference level. It associates part of the range—a pool of numbers—with each tunnel proportional to the tunnel weight. A tunnel with a higher weight is associated with a greater portion of the range—a larger pool—than a tunnel with a lower weight. A tunnel is selected when the random number is in its associated pool of numbers. The random number is more likely, on average, to be in a larger pool, so a tunnel with a higher weight (larger pool) is more likely to be selected than a tunnel with a lower weight (smaller pool).

For example, consider a preference level that has only two tunnels, 1 and 2. Tunnel 1 has a maximum limit of 1000 sessions and Tunnel 2 has a limit of 2000 sessions, resulting in an aggregate total of 3000 sessions. The LAC generates a random number from a pool of 3000 in the range from 0 through 2999. A pool of 1000 numbers, the portion of

the range from 0 through 999, is associated with Tunnel 1. A pool of 2000 numbers, the portion of the range from 1000 through 2999, is associated with Tunnel 2.

- When the generated number is less than 1000, then Tunnel 1 is selected, even though it has a lower weight (1000) than Tunnel 2 (2000).
- When the generated number is 1000 or larger, then Tunnel 2 is selected.

Because the pool of possible generated numbers for Tunnel 2 (2000) is twice that for Tunnel 1 (1000), Tunnel 2, *on average*, is selected twice as often as Tunnel 1.

Destination-Equal Load Balancing

Destination-equal load balancing evaluates tunnels according to the number of sessions to the destination and the number of sessions carried by the tunnel in order to spread the session load equally among all tunnels. The tunnel with a destination that has the lowest session count is considered to have the lightest load. This process operates on tunnels at the highest available preference level and uses the following guidelines:

- When each tunnel goes to a separate destination and only one destination has the lowest session count among all destinations, the LAC selects the tunnel to that destination.
- When each tunnel goes to a separate destination and more than one destination has the same lowest session count, the LAC selects a tunnel at random from among the tunnels to these destinations.
- When more than one tunnel goes to the same destination and that destination has the lowest destination session count, the LAC selects from among these tunnels the one that has the lowest total number of tunnel sessions. If the tunnel session count is the same for all these tunnels, then the LAC selects one of them at random.

Consider the following scenarios to better understand tunnel selection behavior when destination-equal load balancing is enabled.

In Scenario 1, every tunnel has a different valid destination and only the destination session count is evaluated:

- Tunnel 1, preference level 1, 192.168.1.1, destination session count = 200
- Tunnel 2, preference level 1, 192.168.2.2, destination session count = 50
- Tunnel 3, preference level 1, 192.168.3.3, destination session count = 300
- Tunnel 4, preference level 1, 192.168.4.4, destination session count = 100

When the first PPP user tries to connect to the domain, the LAC selects Tunnel 2, because it is at the highest preference level, 1, and has the valid destination, B, with the lowest session count, 50.

When additional PPP users try to connect to the domain, the LAC acts as follows:

1. Tunnel 2 continues to be selected until the session count for 192.168.2.2 equals 100, matching the next lowest session count, 192.168.4.4's in Tunnel 4.
2. When the next subscriber logs in, the LAC randomly selects between Tunnel 2 and Tunnel 4, because their destinations have the same session count, and it is lower than that for the other destinations.
3. Whichever tunnel is selected from this pair, the session count for its destination is now 101. The other tunnel is selected when the next subscriber logs in, because it has the lower destination session count of 100. This raises its destination session count to 101, matching the other tunnel.
4. As subscribers continue to log in, the LAC repeats this process, randomly selecting between Tunnel 2 and Tunnel 4 when their session counts match and then selecting the other tunnel with the next subscriber, until their destination session counts both reach 200, matching Tunnel 1.
5. When the next subscriber logs in, the LAC now randomly selects among Tunnel 1, Tunnel 2, and Tunnel 4, because 192.168.1.1, 192.168.2.2, 192.168.3.3 all have the same session count of 200. The destination session count is raised for the selected tunnel to 201, so for the next subscriber, the LAC randomly selects between the other two tunnels. Now two tunnels have a destination session count of 201, so the LAC selects the remaining tunnel for the next subscriber.
6. As subscribers continue to log in, the LAC repeats this process, randomly selecting among Tunnel 1, Tunnel 2, and Tunnel 4 when their session counts match, randomly selecting between the remaining pair for the next subscriber, and then selecting the remaining tunnel, so the destination session counts for these three tunnels match again. This pattern continues until the destination session count for all three tunnels reaches 300, matching Tunnel 3.
7. Now the destinations for all four tunnels have the same session count. Because there are only four tunnels, the final pattern is established. The LAC first randomly selects among all four tunnels, then the remaining three, then the remaining pair, and finally selects the last tunnel. When the destination session counts are all the same, the LAC starts this pattern again.

In Scenario 2, two tunnels share the same valid destination. The tunnel session count and the destination session count are both evaluated:

- Tunnel 1, preference level 1, tunnel session count = 120, 192.168.1.1, destination session count = 200
- Tunnel 2, preference level 1, tunnel session count = 80, 192.168.1.1, destination session count = 200
- Tunnel 3, preference level 1, 192.168.2.2, destination session count = 300
- Tunnel 4, preference level 2, 192.168.3.3, destination session count = 100

When the first PPP user tries to connect to the domain, the LAC first selects between destinations. The tunnels for both 192.168.1.1 and 192.168.2.2 are at preference level 1. The LAC selects 192.168.1.1, because it has a lower session count (200) than 192.168.2.2 (300). The LAC then has to choose between Tunnel 1 and Tunnel 2 because both go to 192.168.1.1. The LAC evaluates the tunnel session count. Tunnel 2 has a lower count (80) than Tunnel 1 (120), so the LAC selects Tunnel 2 for the first subscriber.

When additional PPP users try to connect to the domain, the LAC acts as follows:

1. Tunnel 2 continues to be selected until its tunnel session count increases to 120, matching Tunnel 1.
2. When the next subscriber logs in, the LAC randomly selects between Tunnel 1 and Tunnel 2, because they have the same tunnel session count. The tunnel session count of the selected tunnel is raised to 121.
3. When the next subscriber logs in, the LAC selects the other tunnel to 192.168.1.1, because it has a lower tunnel session count. From this point, the LAC continues to alternate, first making a random selection between Tunnels 1 and 2 and then selecting the other tunnel, until the destination session count rises to 300, matching the session count for 192.168.2.2 in Tunnel 3. (At this point, the tunnel session count is 150 for both Tunnel 1 and Tunnel 2.)
4. For the next subscriber, the LAC randomly selects among Tunnels 1, 2, and 3.
 - If the LAC selects either Tunnel 1 or Tunnel 2, the 192.168.1.1 session count rises to 301. Consequently the LAC selects Tunnel 3 for the next subscriber because the 192.168.2.2 session count is still 300. At this point, both destinations have the same session count again.
 - If the LAC selects Tunnel 3, the 192.168.2.2 session count rises to 301. For the next subscriber, the LAC randomly selects between Tunnel 1 and Tunnel 2 because they both go to 192.168.1.1. Whichever one the LAC selects, the 192.168.1.1 session count rises to 301. At this point, both destinations have the same session count again.



NOTE: The tunnel session count for Tunnels 1 and 2 is no longer evaluated; the LAC only considers the destination session count for 192.168.1.1 and 192.168.2.2.

This pattern continues for all subsequent subscribers.

In Scenario 3, each tunnel has a different valid destination and only the destination session count is evaluated:

- Tunnel 1, preference level 1, 192.168.1.1, destination session count = 100
- Tunnel 2, preference level 1, 192.168.2.2, destination session count = 100
- Tunnel 3, preference level 1, 192.168.3.3, destination session count = 100
- Tunnel 4, preference level 1, 192.168.4.4, destination session count = 100

When the first PPP user tries to connect to the domain, the LAC determines that the destination session count is the same for all destinations for all four tunnels at the preference level. Consequently, the LAC selects randomly among the four tunnels.

Suppose the LAC selects Tunnel 1 for the first subscriber.

When additional PPP users try to connect to the domain, the LAC acts as follows:

1. The LAC selects randomly among Tunnels 2, 3, and 4, because Destinations 192.168.2.2, 192.168.3.3, and 192.168.4.4 all have the same session count, 100, which is lower than the current session count for 192.168.1.1, 101.
2. Suppose the LAC selects Tunnel 2. For the next subscriber, the LAC randomly selects between Tunnels 3 and 4, because 192.168.3.3 and 192.168.4.4 all have the same session count, 100, which is lower than the current session count of 101 for 192.168.1.1 and 192.168.2.2.
3. Suppose the LAC selects Tunnel 3. For the next subscriber, the LAC selects Tunnel 4, because 192.168.4.4 has a session count of 100, and all the other destinations have a count of 101.
4. Now the destinations for all four tunnels have the same session count. Because there are only four tunnels, the final pattern is established. As subscribers continue to log in, the LAC first randomly selects among all four tunnels, then the remaining three, then the remaining pair, and finally selects the last tunnel. When the destination session counts are all the same, the LAC starts this pattern again.

In Scenario 4, the LAC evaluates both destination session limits and tunnel maximum session limits:

- Tunnel 1, preference level 1, 192.168.1.1, destination session count = 30, tunnel maximum session limit = 200
- Tunnel 2, preference level 1, 192.168.2.2, destination session count = 40, tunnel maximum session limit = 200
- Tunnel 3, preference level 1, 192.168.3.3, destination session count = 300, tunnel maximum session limit = 1000
- Tunnel 4, preference level 2, 192.168.4.4, destination session count = 100

When the first PPP user tries to connect to the domain, the LAC selects Tunnel 1, because 192.168.1.1 has the lowest session count in the preference level.

When additional PPP users try to connect to the domain, the LAC acts as follows:

1. The LAC continues to select Tunnel 1 until the destination session count for 192.168.1.1 equals 40, matching the count for 192.168.2.2 in Tunnel 2.
2. When the next subscriber logs in, the LAC randomly selects between Tunnel 1 and Tunnel 2, because their destinations have the same session count, and it is lower than that for Tunnel 3 (300).
3. Whichever tunnel is selected from this pair, the session count for its destination is now 41. The other tunnel is selected when the next subscriber logs in, because it has

the lower destination session count of 40. This raises its destination session count to 41, matching the other tunnel.

4. As subscribers continue to log in, the LAC repeats this process, randomly selecting between Tunnel 1 and Tunnel 2 when their session counts match and then selecting the other tunnel with the next subscriber, until their destination session counts both reach 200, matching their tunnel maximum session limit of 200. Because both tunnels have reached their maximum session limit, they are not available for selection.
5. As subscribers continue to log in, the LAC selects the remaining tunnel in the preference level, Tunnel 3, until the session count for its destination reaches the maximum session limit for the tunnel, 1000.
6. When the next subscriber logs in, the LAC drops to the next preference level and selects Tunnel 4, because it is the only tunnel at this level.
7. As subscribers continue to log in, the LAC continues to select Tunnel 4, because no maximum session limit is configured for this tunnel. The LAC can subsequently select a tunnel in the higher preference level only when a session is terminated for one of the tunnels at that level, dropping its session count below the maximum limit.

In Scenario 5, one of the destinations is locked:

- Tunnel 1, preference level 1, 192.168.1.1, destination session count = 100, destination locked out
- Tunnel 2, preference level 1, 192.168.2.2, destination session count = 200
- Tunnel 3, preference level 1, 192.168.3.3, destination session count = 250

When the first PPP user tries to connect to the domain, the LAC cannot select Tunnel 1, even though its destination has the lowest session count, because the tunnel is in the destination lockout state. Tunnel 1 cannot be considered until it is out of the locked state. The LAC selects Tunnel 2 because the session count for 192.168.2.2 is lower than for 192.168.3.3.

When additional PPP users try to connect to the domain, what happens next depends on when 192.168.1.1 emerges from the lockout state. For as long as 192.168.1.1 is locked out, the LAC makes the selections as follows:

1. The LAC continues to select Tunnel 2 until the session count for 192.168.2.2 equals 250, matching the count for 192.168.3.3 in Tunnel 3.
2. When the next subscriber logs in, the LAC randomly selects between Tunnel 2 and Tunnel 3, because their destinations have the same session count, 250.
3. Whichever tunnel is selected from this pair, the session count for its destination is now 251. The other tunnel is selected when the next subscriber logs in, because it has the lower destination session count of 250. This raises its destination session count to 251, matching the other tunnel.
4. As subscribers continue to log in, the LAC repeats this process, randomly selecting between Tunnel 2 and Tunnel 3 when their session counts match and then selecting the other tunnel with the next subscriber.

Whenever 192.168.1.1 emerges from the lockout state, the LAC selects Tunnel 1 for the next subscriber because 192.168.1.1 has the lowest session count. The LAC continues to do so until the session count for 192.168.1.1 matches the current session count for either of the other destinations. From that point forward, the LAC alternates making a random selection between tunnels with matching destination session counts and then subsequently selecting the tunnel with the lowest count.

Whenever 192.168.1.1 emerges from the lockout state,

1. The LAC selects Tunnel 1 for the next subscriber because 192.168.1.1 has the lowest session count.
2. The LAC continues to select Tunnel 1 until the session count for 192.168.1.1 matches the current session count for either of the other destinations.
3. From that point forward, the LAC alternates making a random selection between tunnels with matching destination session counts and then subsequently selecting the tunnel with the lowest count.

**Related
Documentation**

- [Configuring the L2TP LAC Tunnel Selection Parameters on page 209](#)
- [Configuring the L2TP Destination Lockout Timeout on page 165](#)
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)
- *Specifying a Tunnel Profile in a Domain Map*

L2TP Session Limits Overview

When an L2TP session request is initiated, the LNS or LAC checks the number of current active sessions against the maximum number of sessions allowed for the chassis, tunnels, a tunnel group, a client (requesting host device), or a group of clients. New session requests are rejected when the configured session limit is reached.

When a session is requested, the LNS checks for session limits in the following order:

chassis > tunnel > tunnel group > session-limit group > client

At each level, the LNS determines whether the current session count is less than the configured limit. When that is true or when no limit is configured, the check passes and the LNS proceeds to check the next level. If at any level the current session count is equal to the configured limit, then the LNS rejects the session request and does not check any other level. Otherwise, the session can be established.

When a session request is rejected for an existing tunnel, a Call-Disconnect-Notify (CDN) message with a result code and error code both set to 4 is returned in response to the incoming-call request (ICRQ). When the rejected request is for a new tunnel, the tunnel is established but the session fails to come up, causing the tunnel to come down because it has no sessions.

The LAC performs the same check, but only for the chassis and tunnel levels. The LAC rejects requests by returning a PPP terminate message to the client.

You can configure session limits for the chassis, all tunnels, a tunnel group, a group of clients, or an individual client. The scenarios that follow describe what happens for different configurations of session limits.

Scenario 1: Chassis Limit

In [Table 14 on page 200](#), the current L2TP session count is 10,000 and the session limit is configured as 10,000 at every level. When a new session is requested, the first check at the chassis level fails, because the current session count matches the configured limit. No further checks are performed at the other levels and the session request is rejected. No new sessions are allowed at any level until the current session count drops below 10,000.

Table 14: Scenario 1, Chassis Limit

Level	Configured Session Limit	Current Session Count Displayed by <code>show services l2tp summary</code> Command	Session Limit Check Result
Chassis	10,000	10,000	Fail
Tunnel A	10,000	10,000	—
Tunnel group B	10,000	10,000	—
Session-limit group	10,000	10,000	—
Client	10,000	10,000	—

Scenario 2: Tunnel Limit

In [Table 15 on page 200](#), the current L2TP session count is 2000. When a new session is requested, the first check at the chassis level passes because the configured limit allows up to 10,000 sessions on the chassis, but only 2000 sessions are currently active. The next check, at the tunnel level, fails, because the current session count matches the configured limit tunnel limit of 2000 for tunnel A.

No further checks are performed at the other levels and the session request is rejected.

Table 15: Scenario 2, Tunnel Limit

Level	Configured Session Limit	Current Session Count Displayed by <code>show services l2tp summary</code> Command	Session Limit Check Result
Chassis	10,000	2000	Pass
Tunnel A	2000	2000	Fail
Tunnel group B	10,000	2000	—
Session-limit group	6000	2000	—
Client	6000	2000	—

No new sessions are allowed on tunnel A until its current session count drops below 2000 and the session check can pass. If that happens, then the other level checks pass in this scenario because their configured limits are greater than their current counts.

The session limit of 2000 applies to all tunnels; that is, each active tunnel has an independent limit of 2000 sessions. The failure of one tunnel has no effect on other tunnels. A session request on any other tunnel passes, as long as the current session count for that tunnel is less than 2000.

Scenario 3: Tunnel Group Limit

In [Table 16 on page 201](#), the current L2TP session count is 2000. When a new session is requested, the first check at the chassis level passes because the configured limit allows up to 10,000 sessions on the chassis, but only 2000 sessions are currently active. The second check, at the tunnel level, also passes for the same reason. The next check, at the tunnel group level for tunnel group B, fails, because the current session count for tunnel group B matches the configured limit tunnel group limit of 2000.

No further checks are performed at the other levels and the session request is rejected.

Table 16: Scenario 3, Tunnel Group Limit

Level	Configured Session Limit	Current Session Count Displayed by show services l2tp summary Command	Session Limit Check Result
Chassis	10,000	2000	Pass
Tunnel A	10,000	2000	Pass
Tunnel group B	2000	2000	Fail
Session-limit group	6000	2000	—
Client	6000	2000	—

No new sessions are allowed on tunnel group B until its current session count drops below 2000 and the session check can pass. If that happens, then the other level checks can pass because their configured limits are greater than their current counts.

For tunnel groups, the session limit is configured on a per-group basis; that is, you cannot specify a single limit that applies to all tunnel groups. The failure of any tunnel group has no effect on other tunnel groups. In this scenario, a session request on any other tunnel group passes, if the current session count for that group is less than its configured session limit.

Scenario 4: Session-Limit Group Limit

In [Table 17 on page 202](#), the current L2TP session count is 6000. When a new session is requested, the check passes for the chassis, tunnel, and tunnel group because the configured limit for each allows up to 10,000 sessions, but only 6000 sessions are

currently active. The check at the session-limit group fails, because the current session count for session-limit group slg1 matches the configured limit of 6000.

No further checks are performed at the remaining level and the session request is rejected.

Table 17: Scenario 4, Session-Limit Group Limit

Level	Configured Session Limit	Current Session Count Displayed by show services l2tp summary Command	Session Limit Check Result
Chassis	10,000	6000	Pass
Tunnel A	10,000	6000	Pass
Tunnel group B	10,000	6000	Pass
Session-Limit group slg1	6000	6000	Fail
Client	8000	2000	—

No new sessions are allowed for any clients in session-limit group slg1 until the group's current session count drops below 6000 and the session check can pass. If that happens, then the remaining level check can pass because its configured limit is greater than its current count.

You can reconfigure a session-limit group by removing or adding clients without affecting any current sessions. The reconfiguration does affect the number of sessions available to be established for the client group.

- If you remove a client, then the number of new sessions that can be established increases by the number of that client's current sessions.
- If you add a client, then the number of new sessions that can be established is reduced by the number of that client's current sessions. The new total of current sessions for existing clients plus the new client can exceed the configured limit for the session-limit group. In this case, no sessions are dropped, but no new sessions can be established until the session count drops below the configured group limit.

To explore this further, consider the following sequence of events:

1. The session-limit group slg1 has two clients, ent1-serviceA with a current session count of 3500 and ent1-serviceB with a current session count of 0. Because group slg1 has a limit of 6000, no more than 2500 sessions can be added for these clients:

$$6000 - 3500 = 2500$$

2. Then 1000 sessions are logged in for client ent1-service B. Now no more than 1500 sessions can be added for these clients:

$$6000 - (3500 + 1000) = 1500$$

- Next, suppose you remove client ent1-serviceA from the session-limit group. The group session capacity increases to 5000 sessions:

$$6000 - 1000 = 5000$$

- Finally, you add a new client, ent1-serviceC, to the session-limit group. This new client currently has 8000 active sessions. In this case, the session-limit group now has 9000 sessions:

$$1000 + 8000 = 9000$$

No sessions are dropped even though the maximum session limit for the group, 6000, is exceeded. No new sessions can be added until the session count drops from 9000 to below 6000.

Scenario 5: Individual Client Limit

In [Table 18 on page 203](#), the session check passes for the chassis, tunnel, and tunnel group because their configured limits are greater than their current session counts. The client, ent1-serviceA, does not belong to a session-limit-group. The limit check fails for the client because its current session count matches the configured limit of 6000.

Table 18: Scenario 5, Individual Client Limit

Level	Configured Session Limit	Current Session Count Displayed by show services l2tp summary Command	Session Limit Check Result
Chassis	10,000	6000	Pass
Tunnel A	10,000	6000	Pass
Tunnel group B	8000	6000	Pass
Client ent1-serviceA	6000	6000	Fail

No new sessions are allowed for this client until its current session count drops below 6000 and the session check can pass. The failure of any independent client has no effect on other clients. In this scenario, a session request for any other independent client passes, if the current session count for that client is less than its configured session limit.

The session limit that you set for an individual client—one that is not part of a session-limit group—applies on a per-tunnel-group basis. Multiple LACs with the same source hostname but different source IP addresses are treated as the same client.

Suppose you have three LACs, A, B, and C. All three have the same source hostname, ce-lac. LAC A and LAC B establish sessions with an LNS through the gateway address associated with tunnel group 1. LAC C establishes sessions through a different gateway associated with tunnel group 2. Because the LACs have the same hostname, the client configuration is the same for all three. However, the client session limit applies differently to the LACs because of the tunnel groups.

Suppose the client session limit is 100. Because LAC A and LAC B both create sessions in tunnel group 1, they must share the client limit. That means that the total number of sessions allowed for LAC A and LAC B combined is 100.

LAC C creates sessions in a different tunnel group, 2. Because the client session limit applies per tunnel group, then LAC C is allowed 100 sessions, regardless of how many sessions LAC A and LAC B have already established.

**Related
Documentation**

- [L2TP Session Limits and Load Balancing for Service Interfaces on page 259](#)
- [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204](#)
- [L2TP for Subscriber Access Overview on page 147](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP LAC on page 173](#)

Limiting the Number of L2TP Sessions Allowed by the LAC or LNS

You can place a limit on the maximum number of L2TP sessions allowed for the chassis, all tunnels, a tunnel group, a group of clients, an individual client, or an individual service interface or aggregated service interface. New session requests are rejected by the LNS or LAC when the configured session limit is reached. Session requests are also rejected when the maximum chassis limit has been reached, even when a configured limit is not exceeded. Configurable session limits provide fine-grained control of the number of sessions that a customer can have while connected over LACs in multiple locations.



NOTE: You cannot set the limit to be more than the default maximum limit for the chassis. For information about the default limit for the different MX Series router models by Junos OS software release, see the *Subscriber Management Scaling Values (XLS)* spreadsheet in the Downloads area at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/index.html.

To limit the number of sessions allowed on a chassis (LAC or LNS):

- Configure the maximum number of sessions.

```
[edit services l2tp]  
user@host# set maximum-sessions number
```

To limit the number of sessions per tunnel for all tunnels (LAC or LNS):

- Configure the maximum number of sessions.

```
[edit services l2tp tunnel ]  
user@host# set maximum-sessions number
```

You cannot set the limit to be more than 65,535 sessions.

To limit the number of sessions for all tunnels in a specific tunnel group (LNS):

- Configure the maximum number of sessions.

```
[edit services l2tp tunnel-group tunnel-group-name]
user@host# set maximum-sessions number
```

To limit the number of sessions that are allowed on an individual service interface:

- Configure the maximum number of sessions.

```
[edit interfaces si-slot/pic/port]
user@host# set l2tp-maximum-session number
```

To limit the number of sessions that are allowed on an individual aggregated service interface:

- Configure the maximum number of sessions.

```
[edit interfaces asinumber]
user@host# set l2tp-maximum-session number
```



NOTE: The configuration applies to all member interfaces; the limit cannot be configured for individual member interfaces of the aggregated service interface.

To limit the number of sessions for a group of clients (LNS):

1. Configure the maximum number of sessions.

```
[edit services l2tp sessions-limit-group limit-group-name]
user@host# set maximum-sessions number
```

2. Associate a client with the session-limit group.

```
[edit access profile profile-name client client-name l2tp]
user@host# set sessions-limit-group limit-group-name
```

To limit the number of sessions for a client that is not a member of a session-limit group (LNS):

- Configure the maximum number of sessions.

```
[edit access profile profile-name client client-name]
user@host# set maximum-sessions number
```



NOTE: Configuring the session limit at any level to be less than the number of sessions that currently exist at that level has no effect on existing sessions. The new limit applies only if the number of sessions drops below the new limit.

Related Documentation

- [L2TP Session Limits Overview on page 199](#)
- [L2TP for Subscriber Access Overview on page 147](#)
- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring L2TP Tunnel Groups](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)
- [Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255](#)
- [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204](#)
- [L2TP Session Limits and Load Balancing for Service Interfaces on page 259](#)

Setting the Format for the Tunnel Name

By default, the name of a tunnel corresponds to the Tunnel-Assignment-Id [82] returned by the AAA server. You can optionally configure the LAC to use more elements in the construction of a tunnel name by including the **assignment-id-format client-server-id** statement at the **[edit services l2tp tunnel]** hierarchy level. This format uses three attributes: Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. These attributes correspond, respectively, to the values configured in the tunnel profile for the LAC (source gateway) name, the tunnel endpoint (remote gateway) address on the LNS, and the tunnel ID.

A consequence of the **client-server-id** format is that the LAC automatically creates a new tunnel when the AAA server returns a different Tunnel-Client-Auth-Id than previously returned.



NOTE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no assignment-id-format assignment-id** statement at the **[edit services l2tp tunnel]** hierarchy level.

To change how the tunnel name is formatted:

- Configure the format.

```
[edit services l2tp tunnel]
user@host# set assignment-id-format client-server-id
```

Related Documentation

- [Configuring an L2TP LAC on page 173](#)

Configuring a Tunnel Profile for Subscriber Access

The tunnel profile specifies a set of attributes to characterize the tunnel. The profile can be applied by a domain map or automatically when the tunnel is created.



NOTE: RADIUS attributes and VSAs can override the values you configured by a tunnel profile in a domain map. In the absence of a domain map, RADIUS can supply all the characteristics of a tunnel. The steps in the following procedure list the corresponding standard RADIUS attribute or VSA that you can configure on your RADIUS server to modify or configure the tunnel profile.

RADIUS-supplied attributes are associated with a tunnel by a tag carried in the attribute, which matches the tunnel identifier. A tag of 0 indicates the tag is not used. If L2TP receives a RADIUS attribute with a tag of 0, the attribute cannot be merged with the tunnel profile configuration corresponding to the subscriber domain because a tunnel profile cannot provide a tunnel tag (tunnel identifier) of 0. Only tags in the range of 1 through 31 are supported.

To configure a tunnel definition for a tunnel profile:

1. Specify the tunnel profile for which you are defining a tunnel. (Tunnel-Group [26-64])

```
[edit access]
user@host# set tunnel-profile profile-name
```

2. Specify an identifier (name) for the L2TP control connection for the tunnel.

```
[edit access tunnel-profile profile-name]
user@host# set tunnel tunnel-id
```

3. Configure the IP address of the local L2TP tunnel endpoint, the LAC. (Tunnel-Client-Endpoint [66])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
user@host# set source-gateway address client-ip-address
```

4. Configure the IP address of the remote L2TP tunnel endpoint, the LNS. (Tunnel-Server-Endpoint [67])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
user@host# set remote-gateway address server-ip-address
```

5. (Optional) Configure the preference level for the tunnel. (Tunnel-Preference [83])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
user@host# set preference number
```

6. (Optional) Configure the hostname of the local client (LAC). (Tunnel-Client-Auth-Id [90])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set source-gateway gateway-name client-name
```

7. (Optional) Configure the hostname of the remote server (LNS). (Tunnel-Server-Auth-Id [91])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set remote-gateway gateway-name server-name
```

8. (Optional) Specify the medium (network) type for the tunnel. (Tunnel-Medium-Type [65])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set medium type
```

9. (Optional) Specify the protocol type for the tunnel. (Tunnel-Type [64])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set type tunnel-type
```

10. (Optional) Configure the assignment ID for the tunnel. (Tunnel-Assignment-Id [82])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set identification name
```

11. (Optional) Configure the maximum number of sessions allowed in the tunnel. (Tunnel-Max-Sessions [26-33])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set max-sessions number
```

12. (Optional) Configure the password for remote server authentication. (Standard RADIUS attribute Tunnel-Password [69] or VSA Tunnel-Password [26-9])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set secret password
```

13. (Optional) Configure the logical system to use for the tunnel.

If you configure a logical system, you must also configure a routing instance.

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set logical-system logical-system-name
```

14. (Optional) Configure the routing instance to use for the tunnel. (Tunnel-Virtual-Router [26-8])

If you configure a routing instance, configuring a logical system is optional.

```
[edit access tunnel-profile profile-name tunnel tunnel-id]  
user@host# set routing-instance routing-instance-name
```

15. (Optional) Enable the LAC to interoperate with Cisco LNS devices.
(Tunnel-Nas-Port-Method [26-30])

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
user@host# set nas-port-method cisco-avp
```

The following example shows a complete configuration for a tunnel profile:

```
tunnel-profile marketing {
  tunnel 1 {
    preference 5;
    remote-gateway {
      address 198.51.100.4;
      gateway-name work;
    }
    source-gateway {
      address 192.0.2.10;
      gateway-name local;
    }
    secret $ABC123;
    logical-system bos-metro-5;
    routing-instance rox-12-32;
    medium ipv4;
    type l2tp;
    identification tunnel_to_work;
    max-sessions 32;
    nas-port-method cisco avp;
  }
}
```

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Domain Mapping Overview](#)
- [LAC Interoperation with Third-Party LNS Devices on page 177](#)

Configuring the L2TP LAC Tunnel Selection Parameters

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from the set of tunnels associated with either the PPP user or the PPP user's domain. You can configure how a tunnel is selected and whether certain information is sent by the LAC to the LNS.

To configure tunnel selection parameters:

1. (Optional) Configure how a tunnel is selected when a connection attempt fails.
See [“Configuring LAC Tunnel Selection Failover Within a Preference Level” on page 210](#).
2. (Optional) Configure how sessions are load-balanced among tunnels.
See [“Configuring Weighted Load Balancing for LAC Tunnel Sessions” on page 211](#).

3. (Optional) Configure sessions to be load-balanced among tunnels within a preference level, by distributing the sessions equally among all tunnels.

See “[Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions](#)” on page 211.

Related Documentation

- [LAC Tunnel Selection Overview on page 183](#)

Configuring LAC Tunnel Selection Failover Within a Preference Level

You can configure how LAC tunnel selection continues in the event of a failure to connect. By default, when the router is unable to connect to a destination at a given preference level, it attempts to connect at the next lower level. You can specify that the router instead attempt to connect to another destination at the same level as the failed attempt.

If all destinations at a preference level are marked as unreachable, the router does not attempt to connect to a destination at that level. It drops to the next lower preference level to select a destination.

If all destinations at all preference levels are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. If the connection fails, the router rejects the PPP user session without attempting to contact the remote router.

For example, suppose there are four tunnels for a domain: A, B, C, and D. All tunnels are considered reachable, and the preference levels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

When the router attempts to connect to the domain, suppose it randomly selects tunnel B from preference 0. If it fails to connect to tunnel B, the router excludes tunnel B for five minutes and attempts to connect to tunnel A. If this attempt also fails, the router drops to preference 1. Then suppose the router selects tunnel C. If it also fails to connect to tunnel C, the router excludes tunnel C for five minutes and attempts to connect to tunnel D.

You configure the preference level used for this tunnel selection method in the tunnel profile or the RADIUS Tunnel-Preference [83] attribute.

To enable tunnel selection failover within a preference level:

- Specify failover within preference.

```
[edit services l2tp]  
user@host# set failover-within-preference
```

Related Documentation

- [LAC Tunnel Selection Overview on page 183](#)
- [Configuring the L2TP LAC Tunnel Selection Parameters on page 209](#)
- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)

- [Filtering RADIUS Attributes and VSAs from RADIUS Messages](#)

Configuring Weighted Load Balancing for LAC Tunnel Sessions

By default, the L2TP LAC selects tunnels for new sessions at random from within the highest available preference level. You can configure the LAC to distribute sessions across tunnels at the highest available preference level by evaluating the weight of each tunnel. This method is known as *weighted load balancing*. The weight of a tunnel is proportional to its maximum session limit and the maximum session limits of the other tunnels at the same preference level. When you configure weighted load balancing, the LAC still selects tunnels randomly within a preference level, but on average the sessions are distributed across tunnels in relationship to the tunnel weights.

To configure weighted load balancing:

- Specify load balancing.

```
[edit services l2tp]  
user@host# set weighted-load-balancing
```

Related Documentation

- [LAC Tunnel Selection Overview on page 183](#)
- [Configuring the L2TP LAC Tunnel Selection Parameters on page 209](#)
- [Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions on page 211](#)

Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions

By default, the L2TP LAC selects tunnels for new sessions at random from within the highest available preference level. Starting in Junos OS Release 15.1, you can configure the LAC to distribute sessions equally across all tunnels at the highest available preference level by evaluating the number of sessions to the destinations and the number of sessions carried by the tunnels. This distribution method is known as *destination-equal load balancing*. The LAC selects the tunnel with the lightest load, according to the following guidelines:

- When each tunnel goes to a separate destination and only one destination has the lowest session count among all destinations, the LAC selects the tunnel to that destination.
- When each tunnel goes to a separate destination and more than one destination has the same lowest session count, the LAC selects a tunnel at random from among the tunnels to these destinations.
- When more than one tunnel goes to the same destination and that destination has the lowest destination session count, the LAC selects from among these tunnels the one that has the lowest total number of tunnel sessions. If the tunnel session count is the same for all these tunnels, then the LAC selects one of them at random.

To configure destination-equal load balancing:

- Specify destination-equal load balancing.
`[edit services l2tp]`
`user@host# set destination-equal-load-balancing`

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, you can configure the LAC to distribute sessions equally across all tunnels at the highest available preference level by evaluating the number of sessions to the destinations and the number of sessions carried by the tunnels.

Related Documentation

- [LAC Tunnel Selection Overview on page 183](#)
- [Configuring the L2TP LAC Tunnel Selection Parameters on page 209](#)
- [Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 211](#)

Enabling the LAC for IPv6 Services

You can configure the LAC to create the IPv6 address family (inet6) when tunneling the subscribers to the LNS. IPv6 firewall filters can then be applied by services on the LAC to subscriber traffic. By default, the LAC requires only family inet to enable forwarding into an IP tunnel. The LAC can apply IPv4 firewall filters to the session. Even when family inet6 is included in the dynamic profile, by default it is not created in order to conserve resources, because it is not needed. Consequently IPv6 firewall filters cannot be applied.

To enable IPv6 address family creation and the application of IPv6 firewall filters:

- Configure enabling.
`[edit services l2tp]`
`user@host# set enable-ipv6-services-for-lac`

You can use the `show services l2tp summary` command to display whether the statement is enabled or disabled.

Related Documentation

- [Configuring an L2TP LAC on page 173](#)

CHAPTER 22

Configuring Use of Subscriber Access Line and Connect Speed Information

- [Subscriber Access Line Information Handling by the LAC and LNS Overview on page 213](#)
- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)
- [Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 227](#)
- [Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228](#)
- [Configuring the Reporting and Processing of Subscriber Access Line Information on page 230](#)
- [Preventing the LAC from Sending Calling Number AVP 22 to the LNS on page 235](#)
- [Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds on page 235](#)

Subscriber Access Line Information Handling by the LAC and LNS Overview

Starting in Junos OS Release 14.1, L2TP supports a set of AVPs that convey information about subscriber access lines from the LAC to the LNS. The information originates from an ANCP access node (DSLAM) and is distributed to the LAC by means of either DSL Forum VSAs in ANCP messages or PPPoE intermediate agent tags included in the PPPoE PADI and PADR messages.

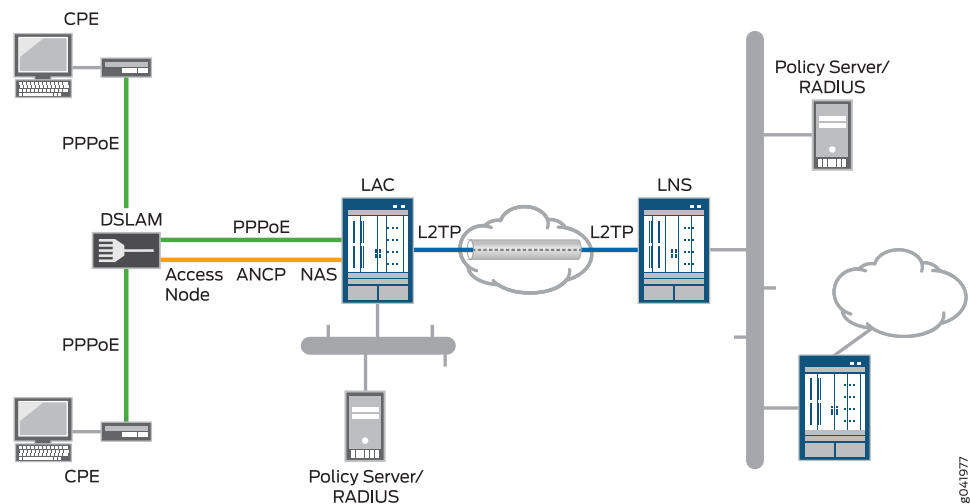
- [Access Line Information Forwarding on page 213](#)
- [Access Line Information AVPs on page 214](#)
- [Connection Speed Updates on the LAC on page 216](#)
- [Connection Speed Updates on the LNS on page 217](#)
- [Interaction Between Global and Per-Destination Configurations on page 217](#)

Access Line Information Forwarding

In the network topology shown in [Figure 13 on page 214](#), when a subscriber initiates a connection through the CPE, the DSLAM relays the subscriber's PPPoE session to the router configured as a LAC. When the router has established the PPPoE session, the LAC initiates an L2TP tunnel to forward the subscriber's encapsulated PPP packets into the provider network.

In parallel to the PPPoE session, an ANCP connection between the DSLAM and the ANCP agent on the router conveys information about the subscriber's local loop as well as the link speeds of the PPPoE sessions on the local loop. The DSLAM sends the router Agent Circuit Id (ACI) and Agent Remote Id (ARI) strings that uniquely identify the DSLAM's receiving interface; this information is encoded in the ANCP Port Up and Port Down messages as Access Line Identifying TLVs. The ANCP messages can also include line attributes such as minimum, maximum, and actual net upstream and downstream data rates in the DSL Line Attributes TLV. The DSLAM can also send the access line attributes in vendor-specific tags that it inserts in the PADI and PADR messages.

Figure 13: Sample L2TP Network Topology



Access Line Information AVPs

L2TP supports the AVPs listed in [Table 19 on page 214](#) to carry this information. The access line information is not required for the L2TP session to be initiated, and the establishment of that session is not delayed waiting for the values to be sent from the DSLAM.

The access line information provided by the AVPs in ICRQ messages is passed on to RADIUS in DSL Forum VSAs. It is not used for shaping the traffic rate on the subscriber access lines.

Table 19: L2TP AVPs That Provide Subscriber Access Line Information

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Actual Data Rate Downstream	130 (26-130)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Actual Data Rate Upstream	129 (26-129)	64-bit unsigned integer; data rate in bits per sec.	ICRQ

Table 19: L2TP AVPs That Provide Subscriber Access Line Information (*continued*)

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Actual Interleaving Delay Downstream	142 (26-142)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Actual Interleaving Delay Upstream	140 (26-140)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Access Loop Encapsulation	144 (26-144)	Three one-octet encodings for data link, encapsulation 1, and encapsulation 2.	ICRQ
Agent Circuit ID	1 (26-1)	2-63 octet string; ACI of the logical access loop port on the DSLAM/access node.	ICRQ
Agent Remote ID	2 (26-2)	2-63 octet statically configured string; uniquely identifies subscriber on the DSLAM (access node).	ICRQ
ANCP Access Line Type	145 (26-145)	One octet encoding for transmission system type, followed by three MBZ (must be zero) octets (total 4 bytes). This value is not supplied in the ICRQ when the access line parameters are sourced from PPPoE-IA, because the ANCP-sourced information may not be immediately available. Starting in Junos OS Release 18.1R1, this AVP is included even when the line type is 0 for OTHER access line types.	ICRQ
Attainable Data Rate Upstream	133 (26-133)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Attainable Data Rate Downstream	134 (26-134)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Connect Speed Update Enable	98 (none)	Value does not matter: presence indicates support for CSUN, CSURQ message types for this session.	ICRQ

Table 19: L2TP AVPs That Provide Subscriber Access Line Information (*continued*)

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Connect Speed Update	97 (none)	Data structure listing remote session id and the current transmit and receive connection speeds in bits per second.	CSUN, CSURQ
IWF Session	254 (26-254)	Four-octet field indicating whether or not the internetworking function has been performed.	ICRQ
Maximum Data Rate Downstream	136 (26-136)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Maximum Data Rate Upstream	135 (26-135)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Maximum Interleaving Delay Downstream	141 (26-141)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Maximum Interleaving Delay Upstream	139 (26-139)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Minimum Data Rate Downstream	132 (26-132)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Downstream Low Power	138 (26-138)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Upstream	131 (26-131)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Upstream Low Power	137 (26-137)	64-bit unsigned integer; data rate in bits per sec.	ICRQ

Connection Speed Updates on the LAC

You can configure the LAC to notify the LNS when the speed of the subscriber connection changes from the values initially communicated to the LNS by AVP 24 (transmit speed) and AVP 38 (receive speed) in Incoming-Call-Connected (ICCN) messages. When configured to do so, the LAC informs the LNS that it can send these updates by including the Connect Speed Update Enable AVP (98) in the ICRQ message when the L2TP session starts up. The absence of the Connect Speed Update Enable AVP (98) in the ICRQ message indicates that the LAC does not send updates for the life of the session.

When the connection speed changes, the DSLAM notifies the ANCP agent. The ANCP agent then notifies the LAC, and the LAC in turn relays this information to the LNS by sending a Connect-Speed-Update-Notification (CSUN) message that includes the updated speeds in a Connect Speed Update AVP (97) for each session. The LAC collects connection speed updates and sends them in a batch to minimize both the performance overhead on the LAC and the amount of traffic generated as a result of these notifications.

The initial speeds in the ICCN messages and updated speeds in CSUN messages are used by CoS to shape the traffic rate for subscriber access lines.

The presence of the Connect Speed Update Enable AVP (98) in the ICRQ message also informs the LNS that the LAC does respond if it receives a Connect-Speed-Update-Request (CSURQ) message from an LNS.



NOTE: The Junos OS does not currently support the sending of CSURQ messages by MX Series routers configured as an LNS. All discussion about CSURQ messages is strictly about how an MX Series LAC responds to a CSURQ that it receives from a third-party LNS.

A third-party LNS can send a CSURQ message at any time during the life of a tunnel to request the current transmit and receive connection speed for one or more L2TP sessions. The LNS includes the remote (relative to the LNS) session IDs in the CSURQ message. If the LAC has previously sent the Connect Speed Update Enable AVP (98) for the requested sessions, then it responds to the CSURQ with a CSUN message that includes the Connect Speed Update AVP (97) for each session. If no changes to connection speeds have occurred by this time, the LAC simply includes the initial connection speed values that were reported in AVP 24 and AVP 38.

When you enable connect speed updates either globally or for a specific LNS, the LAC does not send CSUN messages unless you have also configured the **tx-connect-speed** statement to be either **ancp** or **service-profile**.

Connection Speed Updates on the LNS

Starting in Junos OS Release 17.4R1, an MX Series router configured as an LNS can process subscriber access line information and connection speed updates that it receives from the LAC. The MX Series router cannot send CSURQ messages to solicit updates from the LAC.

The initial speeds in the ICCN messages and updated speeds in CSUN messages are used by CoS to shape the traffic rate for subscriber access lines.

Interaction Between Global and Per-Destination Configurations

You can configure the LAC to forward the access line information in the ICRQ message that it sends to the LNS and you can configure the LNS to receive and process that information. You can configure this globally for all destinations (endpoints) or for a specific destination. The per-destination configuration enables you to limit transmission to an individual LNS or to a set of LNSs or reception from an individual LAC or a set of

LACs. This is useful when you know that some remote gateways do not support this feature or have an incorrect implementation.

Include the **access-line-information** statement at one or both of the following hierarchy levels on the LAC or LNS, respectively, to configure the LAC to forward the access line information in the ICRQ message that it sends to the LNS, or to configure the LNS to receive and process that information:

- **[edit services l2tp]**—Configures forwarding globally for all destinations.
- **[edit services l2tp destination *ip-address*]**—Configures forwarding for a specific destination.

To configure the LAC to send connection speed updates or the LNS to receive and process the updates, include the **connection-speed-update** option with the **access-line-information** statement at the appropriate hierarchy level on the LAC or LNS, respectively.

The global and per-destination settings interact in the following way:

- Access line information—When forwarding by the LAC or processing by the LNS is enabled globally, you cannot disable the global setting for a specific destination.
- Connection speed updates—When forwarding by the LAC or processing by the LNS is enabled globally, you can disable the global setting for a specific destination (LNS or LAC) by specifying **access-line-information** for the destination and omitting **connection-speed-update**.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, an MX Series router configured as an LNS can process subscriber access line information and connection speed updates that it receives from the LAC.
14.1	Starting in Junos OS Release 14.1, L2TP supports a set of AVPs that convey information about subscriber access lines from the LAC to the LNS.

Related Documentation

- [Configuring the Reporting and Processing of Subscriber Access Line Information on page 230](#)
- [L2TP for Subscriber Access Overview on page 147](#)
- [DSL Forum Vendor-Specific Attributes](#)
- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)

Transmission of Tx and Rx Connection Speeds from LAC to LNS

An L2TP access concentrator (LAC) uses Incoming-Call-Connected (ICCN) messages during the establishment of an L2TP tunnel session to send attribute-value pairs (AVP) that convey to the L2TP network server (LNS) the subscriber session's connection speed. AVP 24 includes the transmit (Tx) connect speed and AVP 38 includes the receive (Rx) connect speed.

- The L2TP transmit connect speed is the transmit connect speed in bits per second (bps) of the subscriber's access interface; that is, it represents the speed of the connection downstream from the LAC to the subscriber from the perspective of the LAC.
- The L2TP receive connect speed is the speed in bps of the connection upstream from the subscriber to the LAC, again from the perspective of the LAC. When the receive connect speed is different from the transmit connect speed, AVP 38 is included in the ICCN to convey the receive connect speed.

When the connection speed is the same in both directions, the LNS uses the value in AVP 24 for both transmit and receive connect speeds. In this case, the LAC does not send AVP 38. You can override this default behavior by including the **rx-connect-speed-when-equal** statement, which causes the LAC to send AVP 38 even when the transmit and connect speeds are the same. See [“Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal” on page 227](#).

- The Tx and Rx connect speeds sent in the ICCN message are derived from the method determined by the LAC fallback procedure. Because service activation does not occur until after the ICCN is sent, the LAC always falls back to the next method when **service-profile** is configured as the method. When the service profile is later activated, corresponding speed changes are sent in update messages to the LNS.
- After the L2TP session is established, the Tx and Rx connect speeds can change at any time. When configured to do so, the LAC sends the updated values for each session to the LNS in Connect-Speed-Update-Notification (CSUN) messages. The updated speeds are conveyed in the Connect Speed Update AVP (97).

Methods for Determining the Speed Values Reported to the LNS

The values reported to the LNS can be derived in the following ways:

- You can configure a method globally for the LAC with the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level. You can specify any of the following methods to determine the source for connect speeds:



NOTE: Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release, as described in [Table 20 on page 222](#). The following list includes all historical methods; some of the methods may not be supported in the software release you are using.

- **actual**—The speed is the actual rate of the downstream traffic enforced at the CoS level 3 node based on local traffic control policy. Only the transmit connect speed is available with this method, so the receive transmit speed is determined by the fallback scheme. Use the **actual** method when you need the reported value to be the downstream speed enforced by the local CoS policy. Other methods may vary from this enforced value.

The **actual** method is supported only when the **effective shaping-rate** statement is included at the **[edit chassis]** hierarchy level. The CLI commit check fails if **actual** is configured but the effective shaping rate is not configured.

No commit check is performed when the Tunnel-Tx-Speed-Method VSA (26-94) is set, so a system log message is generated in this situation to remind the user to configure the effective shaping rate.

- **ancp**—The speed is the adjusted ANCP-sourced upstream and downstream value that results from a configured percentage correction to the actual ANCP values. The adjustment is applied on a per-DSL basis to account for ATM encapsulation differences between the BNG and the access-loop and for Layer 1 transport overhead. The initial rate sent to the LNS is the ANCP value reported at the time the ICCN is sent. Any subsequent changes are sent as updates to the LNS in the CSUN message.
- **none**—This option prevents the LAC from sending either AVP 24 or AVP 38 in the ICCN message; consequently no CSUN messages are sent, either. The LNS has to establish its own upstream and downstream policy in the absence of these values. This option overrides the Juniper Networks RADIUS VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163), as well as any other method configured for the connect speed.
- **pppoe-ia-tags**—The speed is derived from the value sent from the DSLAM to the LAC in the Point-to-Point Protocol over Ethernet (PPPoE) intermediate agent (IA) tags. For Ethernet interfaces, the speed is an unadjusted value; for ATM interfaces, the value might be an adjusted value if the tag includes the Encapsulation Overhead attribute (0x90).

This speed value is transmitted when the L2TP session is established. Although the PPPoE IA tag value does not change during a session, the speed reported to the LAC can change. For example, suppose the configured method is **service-profile**. The profile is not activated before the ICCN is sent, and falls back to the PPPoE IA tag, which is sent in the ICCN message. When the service profile is activated later, the service profile rates are sent in an update message (if updates are configured).

- **service-profile**—Depending on your Junos OS release, there are two ways to use service profiles to provide connection speeds. One method uses the speeds from the service profile only in CSUN messages, the other method in ICCN messages.
 - In CSUN messages—The downstream (Tx) speed is derived from the actual CoS that is enforced on the L3 node based on local policy. The upstream (Rx) speed is taken from the configured value in the service profile; no adjustment is made to this value.

By default, service profiles are not activated before the subscriber session is established, so this method falls back to another method for the values sent in

the ICCN. When the profile is later activated, then those rates are sent to the LNS in a CSUN message, if updates are enabled.

- In ICCN messages—Starting in Junos OS Release 18.1R1, you can use a dynamic service profile to provide the connection speeds included in AVP 38 and AVP 24 in the ICCN message when the L2TP session is negotiated. At subscriber login, authd determines whether the service profile name conveyed in the Juniper Networks Activate-Service VSA (26-65) in the RADIUS Access-Accept message matches the service profile name configured with the **service-rate-limiter** statement at the **[edit access]** hierarchy level. If the names match, the speeds are derived either from default values in the service profile or from parameters passed by the VSA. See [“Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds” on page 235](#) for more information about this method.

The **service-profile** method is supported only when the **effective shaping-rate** statement is included at the **[edit chassis]** hierarchy level. The CLI commit check fails when **service-profile** is configured but the effective shaping rate is not configured.

No commit check is performed when the Tunnel-Tx-Speed-Method VSA (26-94) is set, so a system log message is generated in this situation to remind the user to configure the effective shaping rate.



BEST PRACTICE: We recommend that you use only one service profile per subscriber session to affect the downstream shaping rate or report an upstream rate. If more than one dynamic service profile is applied to the subscriber session such that each affects the downstream shaping rate or reports the upstream rate, the values from the most recently applied profile are reported by L2TP. Deactivation of the most recently applied service does not result in L2TP reporting the upstream speed for an existing (active) service profile.

- **static**—This method causes the LAC to derive the speed from the configured static Layer 2 speed. For Ethernet VLANs, this is the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory shaping rate is not configured on the underlying interface, then the actual speed of the underlying physical port is used.
- Starting in Junos OS Release 15.1R1, you can configure speed values directly in the Juniper Networks VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163). These VSAs may be returned in the RADIUS Access-Accept message. If only one of the VSAs is present, the LAC uses a connect speed method to determine the value for the other speed. To use these VSAs, you must configure RADIUS according to your RADIUS server documentation.
- Starting in Junos OS Release 15.1R1, you can configure a method that is conveyed in the Juniper Networks VSA, Tunnel-Tx-Speed-Method (26-94). If configured, this VSA is returned in the RADIUS Access-Accept message for individual subscribers. The VSA value applies globally rather than to a specific tunnel. The method configured in this

VSA specifies the resource that the LAC uses to set the speed. To use this VSA, you must configure RADIUS according to your RADIUS server documentation.

- When the speeds cannot be determined in any other manner, the port speed of the subscriber interface is used.

Table 20 on page 222 lists the available methods by release.



NOTE: Some methods available in VSA 26-94 are not available in the CLI. When one of these methods is received in the VSA, it is translated to a supported method instead of being rejected, or it falls back to another method.

Table 20: Methods for Determining Connect Speeds by Junos OS Release.

Junos OS Release Number	CLI (tx-connect-speed-method)	VSA 26-94 (Tunnel-Tx-Speed-Method)
17.2 and higher	<ul style="list-style-type: none"> • ancp • none • pppoe-ia-tags • service-profile • static (default) 	<ul style="list-style-type: none"> • actual—Translated to service-profile • ancp • CoS—Translated to service-profile • dynamic Layer 2—Translated to static • none • pppoe-ia-tags • service-profile • static
15.1, 16.1, 16.2, 17.1	<ul style="list-style-type: none"> • actual (default) • ancp • none • pppoe-ia-tags 	<ul style="list-style-type: none"> • actual • ancp • CoS—Translated to actual • dynamic Layer 2—Translated to static, which falls back to the port speed of the subscriber access interface • none • pppoe-ia-tags • static—Falls back to the port speed of the subscriber access interface
13.3, 14.1, 14.2	<ul style="list-style-type: none"> • ancp • none • pppoe-ia-tags • static (default) 	n/a



NOTE: Changing the connect speed method in VSA 26-94 or in the CLI configuration has no effect on existing L2TP sessions in which the ICCN has already been sent. All L2TP session negotiations subsequent to the method change use the new setting.

In Junos OS Releases 15.1, 16.1, 16.2, and 17.1 (which support the **actual** method), the speed values in AVP 24 and AVP 38 are typically not greater than the value that is enforced by CoS on the LAC side of the network. Any difference between the speed reported in these AVPs and that enforced by CoS is attributable to differences between the CoS configuration (of the source that is used to enforce a downstream speed) and the Tx connect speed method used to establish these AVPs.

Determining Initial Connect Speeds

Before the LAC can send initial transmit and receive connect speeds in the ICCN message to the LNS, it has to do the following:

1. Select the method it uses to derive the speeds.
2. Determine the speeds.

The LAC selects the method as follows:

1. If the Tunnel-Tx-Speed-Method VSA (26-94) is present, use the method specified by the VSA value.
2. Otherwise, use the method configured in the CLI with the **tx-connect-speed-method** statement.

The LAC determines the initial speed as follows:

1. If the selected method is **none**, the LAC does not include the transmit and receive speeds in the ICCN.
2. For any other selected method, if the values in the Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163) VSAs are nonzero, the LAC sends those values in the ICCN.
3. If the VSA values are zero, use the selected method determined to derive the values to send.

Consider the following examples:

- VSA 26-94 is received with **ancp** configured as the method. The CLI method is configured as **none**. The LAC selects the VSA 26-94 value, the **ancp** method.
VSA 26-162 and VSA 26-163 are received with nonzero values. The LAC sends these VSA values in the ICCN.
- VSA 26-94 is received with **ancp** configured as the method. The CLI method is configured as **none**. The LAC selects the VSA 26-94 value, the **ancp** method.
VSA 26-162 and VSA 26-163 are received with zero values. The LAC uses the **ancp** method to derive the values to send in the ICCN.
- VSA 26-94 is received with **none** configured as the method. The CLI method is configured as **ancp**. The LAC selects the VSA 26-94 value, **none**, and does not send connect speeds in the ICCN.
- VSA 26-94 is not received. The CLI method is configured as **none**. The LAC does not send connect speeds in the ICCN.

Fallback Mechanism for Connect Speed Values

When the LAC has selected a method to derive the connect speeds, it falls back to a different method in any of the following circumstances:

- One or both connect speed values has not been set by the selected method (VSA 26-94 or the CLI).
- The connect speed value is zero.

When one value is available and nonzero but the other is not, only the unset value falls back to a different method. There is no fallback when the selected method is **none**, because this method prevents the LAC from reporting the connect speeds. The fallback procedure can vary by Junos OS release.

Consider the following examples:

- The selected method is ANCP. The ANCP value for the receive speed is found to be zero. The LAC sends the ANCP value for the transmit speed, but the receive value falls back to the PPPoE IA tag method. The LAC sends the IA tag value for the receive speed.
- The selected method is ANCP. The ANCP value for the receive speed is found to be zero. The LAC sends the ANCP value for the transmit speed, but the receive value falls back to the PPPoE IA tag method. The IA tag value for the receive speed is also found to be zero, so it falls back to the static Layer 2 method. This is available, so the LAC sends the static Layer 2 value for the receive speed.
- The selected method is service profile. The service profile is not activated before the ICCN is sent, so the LAC falls back to the ANCP method. Both transmit and receive ANCP values are available and nonzero, so the LAC sends these values in the ICCN.

The service profile is activated by a Change of Authorization (CoA) at some later time for the session. If updates are enabled, the LAC sends the service profile values to the LNS in a CSUN message. If updates are not enabled, the service profile values are not reported to the LNS.

Note that updates require the method to be configured in the CLI. Consequently, VSA 26-94 must not be configured or received so that the service profile method is selected from the CLI configuration.

Starting in Junos OS Release 17.2R1, the LAC fallback procedure is as described in [Table 21 on page 224](#).

Table 21: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Release 17.2 and Higher)

Method	Transmit and Receive Speed Not Set	Transmit Speed Not Set	Receive Speed Not Set
None	No fallback.	No fallback.	No fallback.
Service profile	Both fall back to ANCP method.	Transmit speed falls back to ANCP method.	Receive speed falls back to ANCP method.

Table 21: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Release 17.2 and Higher) (*continued*)

Method	Transmit and Receive Speed Not Set	Transmit Speed Not Set	Receive Speed Not Set
ANCP	Both fall back to PPPoE IA tags method.	Transmit speed falls back to PPPoE IA tags method.	Receive speed falls back to PPPoE IA tags method.
PPPoE IA tags	Both fall back to static Layer 2 method.	Transmit speed falls back to static Layer 2 method.	Receive speed falls back to static Layer 2 method.
Static Layer 2	Both fall back to port speed.	Transmit speed falls back to port speed.	Receive speed falls back to transmit speed.

Starting in Junos OS Release 15.1R1, the LAC fallback procedure is as described in [Table 22 on page 226](#).

Table 22: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Releases 15.1, 16.1, 16.2, 17.1)

Method	Transmit and Receive Speed Not Set	Transmit Speed Not Set	Receive Speed Not Set
None	No fallback.	No fallback.	No fallback.
Actual	Both fall back to ANCP method.	Transmit speed falls back to ANCP method.	Receive speed falls back to ANCP method.
ANCP	Both fall back to PPPoE IA tags method.	If PPPoE IA tags are available for both, then both fall back to PPPoE IA tags method. Otherwise, transmit speed falls back to PPPoE IA tags method.	If PPPoE IA tags are available for both, then both fall back to PPPoE IA tags method. Otherwise, receive speed falls back to PPPoE IA tags method.
PPPoE IA tags	Both fall back to port speed.	Transmit speed falls back to port speed.	Receive speed falls back to port speed.

Starting in Junos OS Release 13.3R1, the LAC fallback procedure is as described in [Table 23 on page 226](#).

Table 23: LAC Fallback Procedure When a Connect Speed Value is Not Set (Junos OS Releases 13.3, 14.1, 14.2)

Method	Transmit and Receive Speed Not Set	Transmit Speed Not Set	Receive Speed Not Set
None	No fallback.	No fallback.	No fallback.
ANCP	Both fall back to PPPoE IA tags method.	If PPPoE IA tags are available for both, then both fall back to PPPoE IA tags method. Otherwise, transmit speed falls back to PPPoE IA tags method.	If PPPoE IA tags are available for both, then both fall back to PPPoE IA tags method. Otherwise, receive speed falls back to PPPoE IA tags method.
PPPoE IA tags	Both fall back to static Layer 2 method.	Transmit speed falls back to static Layer 2 method.	Receive speed falls back to static Layer 2 method.
Static Layer 2	Both fall back to port speed.	Transmit speed falls back to port speed.	Receive speed falls back to transmit speed.



NOTE: For both Gigabit Ethernet (ge) and 10-Gigabit Ethernet (xe) interfaces, the port speed value is set to 1,000,000,000. For aggregated Ethernet (ae) interfaces, the port speed value is set to 0. The port speed value for all these interface types is reported in both AVP 24 and AVP 38.

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, you can use a dynamic service profile to provide the connection speeds included in AVP 38 and AVP 24 in the ICCN message when the L2TP session is negotiated.
17.2R1	Starting in Junos OS Release 17.2R1, the LAC fallback procedure is as described in Table 21 on page 224 .
17.2R1	Starting in Junos OS Release 15.1R1, the LAC fallback procedure is as described in Table 22 on page 226 .
17.2R1	Starting in Junos OS Release 13.3R1, the LAC fallback procedure is as described in Table 23 on page 226 .
15.1R1	Starting in Junos OS Release 15.1R1, you can configure speed values directly in the Juniper Networks VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163).
15.1R1	Starting in Junos OS Release 15.1R1, you can configure a method that is conveyed in the Juniper Networks VSA, Tunnel-Tx-Speed-Method (26-94).
13.3R1	Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release, as described in Table 20 on page 222 .

Related Documentation

- [Juniper Networks VSAs Supported by the AAA Service Framework](#)
- [Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 227](#)
- [Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228](#)
- [Configuring an L2TP LAC on page 173](#)

Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal

The L2TP Rx Connect Speed (in bits per second) AVP, which is represented by AVP 38, is included in the ICCN message when the receive connect speed is different from the transmit connect speed. By default, when the connection speed is the same in both directions, AVP 38 is not sent; the LNS uses the value in AVP 24 for both transmit and receive connect speeds.

AVP 38 is generated when the receive connect speed of the access interface is set equal to the calculated transmit connect speed by issuing the **rx-connect-speed-when-equal** statement at the **[edit services l2tp]** hierarchy level. In this scenario, the LAC transmits the same value for transmit and receive connect speeds that are sent to the LNS through the AVP 24 and AVP 38 in the ICCN message.

To configure the sending of AVP 38 when the connection speeds are the same in both the downstream and upstream directions:

- Configure the transmission of the receive connect speed, AVP 38, when the receive connect speed is set equal to the calculated transmit connect speed.

```
[edit services l2tp]
user@host# set rx-connect-speed-when-equal
```

**Related
Documentation**

- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)
- [Configuring an L2TP LAC on page 173](#)
- [rx-connect-speed-when-equal on page 607](#)

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS

The LAC connection speeds are determined in one of several ways:

- The Juniper Networks VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163).
- The Juniper Networks VSA, Tunnel-Tx-Speed-Method (26-94).
- The CLI configuration.
- The port speed of the subscriber access interface.

You can include the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level to configure a method that specifies the resource that the LAC uses for setting these speeds when the Juniper Networks VSAs are not returned for the subscriber.

Starting in Junos OS Release 17.2R1, when you enable connect speed updates for the LAC you must include the **tx-connect-speed-method** statement. You also must specify either **ancp** or **service-profile** as the method; otherwise, the LAC does not send CSUN messages.

Changing the connect speed method in the CLI configuration or in VSA 26-94 has no effect on existing L2TP sessions in which the ICCN has already been sent. All L2TP session negotiations subsequent to the method change use the new setting.



NOTE: Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release. The following procedure lists all historical methods; some of the methods may not be supported in the software release you are using. See [“Transmission of Tx and Rx Connection Speeds from LAC to LNS” on page 219](#) for a table of support by release.

To set the method for calculating the transmit connect speed:

- (Optional) Configure the LAC to use the class-of-service effective shaping rates.

```
[edit services l2tp]
user@host# set tx-connect-speed-method actual
```




NOTE: This method requires that the effective shaping rate statement is configured at the [edit chassis] hierarchy level. If it is not, then committing this method fails. However, if the method is received from RADIUS in VSA 26-94, a system log message is generated instead, because no commit check is performed in this case.

- (Optional) Configure the LAC to use the values derived from the ANCP value configured on the PPPoE interface underlying the subscriber interface.

```
[edit services l2tp]
user@host# set tx-connect-speed-method ancp
```

- (Optional) Configure the LAC to use the values provided in the PPPoE IA tags received from the DSLAM.

```
[edit services l2tp]
user@host# set tx-connect-speed-method pppoe-ia-tags
```

In this case, the value of Actual-Data-Rate-Downstream (VSA 26-129) is used for AVP 24. The value of Actual-Data-Rate-Upstream (VSA 26-130) is used for AVP 38 and is sent only when the VSA values differ.



NOTE: This speed derived from the IA tags does not apply to subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved.

- (Optional) Configure the LAC to use the following:
 - Downstream (Tx) speed: The actual CoS rate that is enforced on the level 3 node based on local policy
 - Upstream (Rx) speed: The value configured in the dynamic service profile.
- 1. Specify the **service-profile** method.

```
[edit services l2tp]
user@host# set tx-connect-speed-method service-profile
```

2. In the dynamic service profile, configure the ingress shaping rate from CoS to be used by the LAC to report to the LNS as the Rx connect speed.

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
logical-unit-number]
user@host# set report-ingress-shaping-rate bps
```



NOTE: The **service-profile** method requires that the effective shaping rate statement is configured at the [edit chassis] hierarchy level. If it is not, the commit check fails. However, if the **service-profile** method is received from RADIUS in VSA 26-94, a system log message is generated instead, because no commit check is performed in this case.



NOTE: For another method to use service profiles to provide the connection speeds, see [“Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds” on page 235](#).

- (Optional) Configure the LAC to use the underlying interface’s recommended (advisory) downstream shaping rate for AVP 24 and recommended upstream shaping rate for AVP 38. This is also referred to as the static Layer 2 shaping rate.

```
[edit services l2tp]
user@host# set tx-connect-speed-method static
```

You configure the advisory rates under the PPPoE logical interface underlying the subscriber interface with the **advisory-options** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. If the advisory speed is not configured, then the actual port speed is used. For ge and xe interfaces, the speed value is set to 10,000,000 and for ae interfaces, the speed value is set to 0 and sent in both AVP 24 and AVP 38

- (Optional) Configure the LAC to disable sending AVP 24 and AVP 38.



NOTE: This option prevents the LAC from sending either AVP 24 or AVP 38 in the ICCN messages. This option also overrides the Juniper Networks RADIUS VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163).

```
[edit services l2tp]
user@host# set tx-connect-speed-method none
```

Release History Table

Release	Description
17.2R1	Starting in Junos OS Release 17.2R1, when you enable connect speed updates for the LAC you must include the tx-connect-speed-method statement.
13.3R1	Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release.

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)
- [Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds on page 235](#)

Configuring the Reporting and Processing of Subscriber Access Line Information

The L2TP AVP extensions defined in RFC 5515, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extension*, enable the LAC to report to the LNS characteristics of the subscriber’s access line, such as identification attributes, line

type, connection speed, various data rates, and so on. The LAC receives the access line information when the subscriber's CPE initiates a connection request, and forwards the available information in various AVPs included in ICRQ messages to the LNS. The LAC can also signal to the LNS that it is capable of sending updates to the subscriber connection speeds; these are conveyed by the Connect Speed Update AVP (97) in the CSUN message.

Starting in Junos OS Release 17.4R1, RFC 5515 AVP extensions are also supported on the LNS. Consequently, you can configure the LNS to process subscriber access line information and connection speed updates that it receives from the LAC.



NOTE: Subscriber access line information conveyed by AVPs in ICRQ messages is passed to RADIUS in DSL Forum VSA AVPs. Initial and updated connection speeds conveyed in ICCN and CSUN messages can be used by CoS to adjust traffic rates for the subscriber lines.

By default, neither the access line information forwarding or connection speed update capability are enabled on the LAC. You must configure the capabilities for all LNS endpoints or for a specific LNS endpoint. The per-destination configuration applies to all tunnels with that destination IP address. You might want to use a per-destination configuration when you know that only certain endpoints support or correctly implement this feature.

Similarly, processing of this information by the LNS is not enabled by default. You can enable processing for information received from all LAC endpoints or for specific LAC endpoints. The per-destination configuration applies to all tunnels with that destination IP address.



NOTE: The CLI statements are the same for both the LAC and LNS; the difference is that you include the statements in the LAC configuration or the LNS configuration.

To configure the LAC to send information about subscriber access lines to the LNS, or to configure the LNS to process this information received from the LAC:

- Configure the capability globally for all endpoints.

```
[edit services l2tp]
user@host# set access-line-information
```

- Configure the capability for a specific endpoint.

```
[edit services l2tp destination address ip-address]
user@host# set access-line-information
```



BEST PRACTICE: Do not configure the `connection-speed-update` option on the LAC when the LNS does not support connection speed changes. This might be an LNS that is not configured to process the updates or a

noncompliant, third-party LNS. Configuring the LAC option for such an LNS generates additional control messages that are ignored.

.....

To configure the LAC to also send updates to the LNS about changes in connection speed, or to configure the LNS to process speed updates received from the LAC:

- Include the update option when you configure the capability.

```
[edit services l2tp]
user@host# set access-line-information connection-speed-update
```

or

```
[edit services l2tp destination address ip-address]
user@host# set access-line-information connection-speed-update
```

- When you configure the LAC to send updates, you must also configure the method by which the connect speed values are derived. The method specifies the source of the update values. On the LNS, the derivation method is not relevant and cannot be configured.

```
[edit services l2tp]
user@host# set tx-connect-speed-method method
```

Consider the following examples:

- The following configuration specifies that for all tunnels with an endpoint address of 192.0.2.2, the LAC reports access line characteristics sourced from the ANCP agent or the PPPoE intermediate agent (in that order) to the LNS in the ICRQ message. The Connect Speed Update Enable AVP (98) is not included in the ICRQ; consequently no CSUN messages are sent to the LNS to report speed changes in the subscriber access lines reported by the ANCP agent. The LAC ignores any CSURQ messages that it receives from the LNS; this can be only a third-party LNS, because the sending of CSURQ messages is not supported on MX Series routers configured as an LNS.

```
[edit services l2tp destination address 192.0.2.2]
user@host# set access-line-information
```

- The following configuration specifies that for all tunnels with an endpoint address of 203.0.113.23, the LAC reports access line characteristics sourced from the ANCP agent or the PPPoE intermediate agent (in that order) to the LNS in the ICRQ message. The Connect Speed Update Enable AVP (98) is included in the ICRQ; CSUN messages are sent to the LNS to report speed changes in the subscriber access lines reported by the ANCP agent. The LAC accepts any CSURQ messages that it receives from the LNS and responds with a CSUN message; this can be only a third-party LNS, because the sending of CSURQ messages is not supported on MX Series routers configured as an LNS.

```
[edit services l2tp]
user@host# set destination address 203.0.113.23 access-line-information
connection-speed-update
user@host# set tx-connect-speed-method ancp
```

When access line information forwarding is enabled globally, you cannot disable it for a specific destination. However, when connection speed updates are enabled globally, you can disable updates for a specific destination.

- The following configuration specifies that both forwarding of access line characteristics and connection speed updates are enabled for all destinations. For destination 198.51.100.2, the global updates configuration is overridden by repeating the access line configuration for, and omitting the connection speed updates for, that destination.

```
[edit services l2tp]
user@host# set access-line-information connection-speed-update
user@host# set tx-connect-speed-method ancp
[edit services l2tp destination address 198.51.100.2]
user@host# set access-line-information
```

The **show services l2tp summary** command displays the configuration that applies to all destinations. The following sample output confirms the global configuration in this example:

```
user@host> show services l2tp summary
Failover within a preference level is Disabled
Weighted load balancing is Disabled
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 assignment id format is assignment-id
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
Access Line Information is Enabled, Speed Updates is Enabled
Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

The **show services l2tp destination detail** command displays the configuration for each destination individually. The following sample output verifies that connection speed updates are disabled for 198.51.100.2:

```
user@host> show services l2tp destination detail
Local name: 1
Remote IP: 198.51.100.2
Tunnels: 1, Sessions: 1
State: Enabled
Local IP: 203.0.113.2
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Access Line Information: Enabled, Speed Updates: Disabled
...
```

- In this example, the forwarding of access line characteristics is enabled for all destinations, but connection speed updates are enabled for only one destination, 198.51.100.21.

```
[edit services l2tp]
user@host# set access-line-information
[edit services l2tp destination address 198.51.100.21]
user@host# set access-line-information connection-speed-update
user@host# up
user@host# set tx-connect-speed-method ancp
```

The following sample output confirms that connection speed updates are disabled globally:

```
user@host> show services l2tp summary
Failover within a preference level is Disabled
Weighted load balancing is Disabled
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 assignment id format is assignment-id
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
Access Line Information is Enabled, Speed Updates is Disabled
Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

The following sample output confirms that connection speed updates are enabled for destination 198.51.100.21:

```
user@host> show services l2tp destination detail
Local name: 1
Remote IP: 198.51.100.21
Tunnels: 1, Sessions: 1
State: Enabled
Local IP: 203.0.113.3
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Access Line Information: Enabled, Speed Updates: Enabled
...
```

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, RFC 5515 AVP extensions are also supported on the LNS.

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Subscriber Access Line Information Handling by the LAC and LNS Overview on page 213](#)

Preventing the LAC from Sending Calling Number AVP 22 to the LNS

Calling Number AVP 22 typically identifies the interface that is connected to the customer in the access network. When RADIUS includes the Calling-Station-Id in the Access-Accept message, that value is used for the Calling Number AVP. Otherwise, the underlying interface (for example, the S-VLAN IFL) on which the PPPoE session is established is used for the Calling Number AVP value.

By default, the LAC includes this AVP in the incoming-call request (ICRQ) packets that it sends to the LNS. However, you may wish to hide your network access interface information. To do so, you can configure the tunnel so that the LAC does not send the Calling Number AVP to the LNS.

To disable sending the Calling Number AVP:

- Configure disabling.

```
[edit services l2tp]
user@host# set disable-calling-number-avp
```

Related Documentation • [LAC Tunnel Selection Overview on page 183](#)

Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds

When an L2TP session is negotiated, the LAC sends to the LNS an ICCN message that includes values for the Rx connection speed (in AVP 38) and Tx connection speed (in AVP 24) at the LAC. The LAC uses values from the best source available at the time of negotiation. If multiple sources are available, the selection is made based on preference hierarchy of the sources. The source is either RADIUS, ANCP, or PPPoE-IA tags.

By default, the LAC cannot use a service profile received in a RADIUS Access-Accept message as the source, because the profile is not applied until the network family is activated, which occurs after the session negotiation completes. However, if the LNS supports *RFC 5515, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*, the LAC can send a connection speed update to the LNS with values from the service profile.

Starting in Junos OS Release 18.1R1, you can use a dynamic service profile to provide the connection speeds included in AVP 38 and AVP 24 when the L2TP session is negotiated. At subscriber login, authd determines whether the configured service profile name matches the profile name conveyed in the Juniper Networks Activate-Service VSA (26-65) in the RADIUS Access-Accept message. If the names match, the speeds are derived either from default values in the service profile or from parameters passed by the VSA.

This processing by authd to establish the connection speeds takes place only at subscriber login. It does not occur in response to reauthentication or CoA requests.



NOTE: For this feature to work, you must also use the `tx-connect-speed-method` statement at the `[edit services l2tp]` hierarchy level to set the method to `service-profile`. You must also configure the `effective-shaping-rate` statement at the `[edit chassis]` hierarchy level.

You can define the rates directly in the service profile as default values for user-defined variables. Alternatively, you can configure the rates to be passed by RADIUS in VSA 26-65. In either case, the first value is taken as the receive speed (the upstream rate from the subscriber to the LAC) and the second value is taken as the transmit speed (the downstream rate from the LAC to the subscriber). The VSA might be configured to pass more than two parameters, but only the first two parameters matter for the service rate-limiting function.

The rate values are specified in the profile or VSA 26-65 in Kbps, but the L2TP AVP format requires rate values in bps. When you enable this feature, default multipliers automatically convert the rates from Kbps to bps. You can also configure the multiplier options to adjust the rates up or down. The adjusted values are equivalent to the Juniper Networks RADIUS VSAs, Rx-Connect-Speed (26-163) and Tx-Connect-Speed (26-162). These values are stored as such in the session database. Because the values are available in the SDB before the L2TP connection is negotiated, the LAC includes them in the ICCN message as AVP 38 and AVP 24. They are treated as RADIUS-sourced values and consequently have the highest precedence.



NOTE: A parameter value of zero signifies that the rate is not set. For example, if VSA 26-65 returns `service-profile-name(0, 0)`, then no value is set in the SDB for Rx or Tx.

Another circumstance that causes no values to be set in the SDB is if VSA 26-65 does not pass any parameters and you failed to set default values in the service profile. In this case, there are no values for authd to derive and so nothing to place in the SDB for Rx or Tx.

If the service used to establish the rate limiters is deactivated or deleted, authd then clears those rate limiter values from the subscriber session. If the service is reactivated, authd does not reinstate the rate limiters.

To configure LAC connection speeds to be derived at login from a dynamic service profile and to optionally adjust the rates:

1. Specify the dynamic service profile that supplies the connection speeds.

[edit access]

```
user@host# set service-rate-limiter service-name service-profile-name
```

2. (Optional) Configure a value that is multiplied with the Rx connect speed specified in the service profile.

[edit access]


```
user@host# set service-rate-limiter rx-multiplier rx-multiplier
```

3. (Optional) Configure a value that is multiplied with the Tx connect speed specified in the service profile.

```
[edit access]
user@host# set service-rate-limiter tx-multiplier multiplier
```

4. Set the method for determining the connection speed.

```
[edit services l2tp]
user@host# set tx-connect-speed-method service-profile
```

5. Enable the reporting of the actual downstream rate in RADIUS accounting messages.

```
[edit chassis]
user@host# set effective-shaping-rate
```

For example, suppose you configure a dynamic service policy, l2tp-service. The policy includes user-defined variables, upstream and downstream, with default values, respectively, of 20,000 Kbps and 30,000 Kbps. The upstream variable is used for the input (ingress) filter and downstream variable is used for the output (egress) filter.

```
[edit dynamic-profiles l2tp-service]
user@host# set variables upstream default-value 20000
user@host# set variables downstream default-value 30000
user@host# set variables aggregate default-value 50000
user@host# interfaces pp0 "$junos-interface-unit" family inet filter input "$upstream"
user@host# interfaces pp0 "$junos-interface-unit" family inet filter output "$downstream"
```

Then you configure the following service rate limiter, which specifies that when a service policy named l2tp-service is returned, the Rx value in the policy, or passed by the VSA, is multiplied by 1005. The Tx value is multiplied by 1003.

```
[edit access]
user@host# set service-rate-limiter service-name l2tp-service
user@host# set service-rate-limiter rx-multiplier 1005
user@host# set service-rate-limiter tx-multiplier 1003
```

Suppose a subscriber logs in and the Access-Accept message from the RADIUS server includes the Activate-Service VSA, 26-55, specifying l2tp-service. What happens next depends on the parameters passed by the VSA.

- The VSA includes "l2tp-service" with no parameters. The following values are stored in the SDB:
 - Rx is the default value in the policy multiplied by the configured multiplier:
20000 Kbps x 1005 = 20,100,000 bps.
 - Tx is the default value in the policy multiplied by the configured multiplier:
30000 Kbps x 1003 = 30,090,000 bps.
- The VSA includes "l2tp-service(10000, 15000)". The following values are stored in the SDB:

- Rx is the first parameter passed by the VSA multiplied by the configured multiplier:
 $10000 \text{ Kbps} \times 1005 = 10,050,000 \text{ bps}$.
- Tx is the second parameter passed by the VSA multiplied by the configured multiplier:
 $15000 \text{ Kbps} \times 1003 = 15,045,000 \text{ bps}$.
- The VSA includes "l2tp-service(10000)". The following values are stored in the SDB:
 - Rx is the first (and only) parameter passed by the VSA multiplied by the configured multiplier:
 $10000 \text{ Kbps} \times 1005 = 10,050,000 \text{ bps}$.
 - Because the VSA does not pass a second parameter, Tx is the default value in the policy multiplied by the configured multiplier:
 $30000 \text{ Kbps} \times 1003 = 30,090,000 \text{ bps}$.
- The VSA includes "l2tp-service(10000, 0)". The following values are stored in the SDB:
 - Rx is the first parameter passed by the VSA multiplied by the configured multiplier:
 $10000 \text{ Kbps} \times 1005 = 10,050,000 \text{ bps}$.
 - Because the second parameter passed is zero, and zero means that the rate is not set, no value is stored in the SDB for Tx.
- The VSA includes "l2tp-service(0, 0)". The following values are stored in the SDB:
 - Because a passed value of zero means that the rate is not set, no value is stored in the SDB for either Rx or Tx.
- The VSA includes "l2tp-service(10000, 15000, 4000000)". The following values are stored in the SDB:
 - Rx is the first parameter passed by the VSA multiplied by the configured multiplier:
 $10000 \text{ Kbps} \times 1005 = 10,050,000 \text{ bps}$.
 - Tx is the second parameter passed by the VSA multiplied by the configured multiplier:
 $15000 \text{ Kbps} \times 1003 = 15,045,000 \text{ bps}$.
-

Related Documentation

- [Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228](#)
- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)
- [Configuring an L2TP LAC on page 173](#)

CHAPTER 23

Configuring L2TP LNS Inline Service Interfaces

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)
- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)
- [Configuring a AAA Local Access Profile on the LNS on page 248](#)
- [Configuring an Address-Assignment Pool for L2TP LNS with Inline Services on page 249](#)
- [Configuring the L2TP LNS Peer Interface on page 251](#)
- [Enabling Inline Service Interfaces on page 251](#)
- [Configuring an Inline Service Interface for L2TP LNS on page 252](#)
- [Configuring Options for the LNS Inline Services Logical Interface on page 253](#)
- [LNS 1:1 Stateful Redundancy Overview on page 254](#)
- [Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255](#)
- [Verifying LNS Aggregated Inline Service Interface 1:1 Redundancy on page 257](#)
- [L2TP Session Limits and Load Balancing for Service Interfaces on page 259](#)
- [Example: Configuring an L2TP LNS on page 263](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)
- [Applying Services to an L2TP Session Without Using RADIUS on page 277](#)
- [Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 284](#)
- [Configuring a Dynamic Profile for Dynamic LNS Sessions on page 285](#)

Configuring an L2TP LNS with Inline Service Interfaces

The L2TP LNS feature license must be installed before you begin the configuration. Otherwise, a warning message is displayed when the configuration is committed.

To configure an L2TP LNS with inline service interfaces:

1. (Optional) Configure a user group profile that defines the PPP configuration for tunnel subscribers.

See [“Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile” on page 245.](#)

2. (Optional) Configure PPP attributes for subscribers on inline service interfaces.

See [“Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface” on page 242.](#)

3. Configure inline IP reassembly.

See [“Configuring IP Inline Reassembly for L2TP” on page 290.](#)

4. Configure an L2TP access profile that defines the L2TP parameters for each LNS client (LAC).

See [“Configuring an L2TP Access Profile on the LNS” on page 246.](#)

5. (Optional) Configure a AAA access profile to override the access profile configured under the routing instance.

See [“Configuring a AAA Local Access Profile on the LNS” on page 248.](#)

6. Configure a pool of addresses to be dynamically assigned to tunneled PPP subscribers.

See [“Configuring an Address-Assignment Pool for L2TP LNS with Inline Services” on page 249.](#)

7. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address.

See [“Configuring the L2TP LNS Peer Interface” on page 251.](#)

8. Enable inline service interfaces on an MPC.

See [“Enabling Inline Service Interfaces” on page 251.](#)

9. Configure a service interface.

See [“Configuring an Inline Service Interface for L2TP LNS” on page 252.](#)

10. Configure options for each inline service logical interface.

See [“Configuring Options for the LNS Inline Services Logical Interface” on page 253.](#)

11. (Optional) Configure an aggregated inline service interface and 1:1 stateful redundancy.

See [“Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces” on page 255](#)

12. Configure the L2TP tunnel group.

See [“Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces” on page 274](#).

13. (Optional) Configure a dynamic profile that dynamically creates L2TP logical interfaces.

See [“Configuring a Dynamic Profile for Dynamic LNS Sessions” on page 285](#).

14. (Optional) Configure a service interface pool for dynamic LNS sessions.

See [“Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions” on page 284](#).

15. (Optional) Specify how many times L2TP retransmits unacknowledged control messages.

See [“Configuring Retransmission Attributes for L2TP Control Messages” on page 171](#).

16. (Optional) Specify how long a tunnel can remain idle before being torn down.

See [“Setting the L2TP Tunnel Idle Timeout” on page 163](#).

17. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

See [“Setting the L2TP Receive Window Size” on page 163](#).

18. (Optional) Specify how long the L2TP retains information about terminated dynamic tunnels, sessions, and destinations.

See [“Setting the L2TP Destruct Timeout” on page 164](#).

19. (Optional) Configure the L2TP destination lockout timeout.

See [“Configuring the L2TP Destination Lockout Timeout” on page 165](#).

20. (Optional) Configure L2TP tunnel switching.

See [“Configuring L2TP Tunnel Switching” on page 161](#).

21. (Optional) Prevent the creation of new sessions, destinations, or tunnels for L2TP.

See [“Configuring L2TP Drain” on page 166](#).

22. (Optional) Configure whether the L2TP failover protocol is negotiated or the silent failover method is used for resynchronization.

See [“Configuring the L2TP Peer Resynchronization Method” on page 295](#).

23. (Optional) Enable SNMP statistics counters.

See [“Enabling Tunnel and Global Counters for SNMP Statistics Collection” on page 303](#).

24. (Optional) Configure trace options for troubleshooting the configuration.

See [“Tracing L2TP Operations for Subscriber Access” on page 371](#).

You also need to configure CoS for LNS sessions. For more information, see *Configuring Dynamic CoS for an L2TP LNS Inline Service*.

Related Documentation

- [L2TP for Subscriber Access Overview on page 147](#)
- *Junos OS Feature Licenses*
- *Software Feature Licenses*

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

You can configure PPP attributes that are applied by the LNS on the inline service (si) interface to the PPP subscribers tunneled from the LAC. Because you are configuring the attributes per interface rather than with a user group profile, the attributes for subscribers can be varied with a finer granularity. This configuration matches that used for terminated PPPoE subscribers.

To configure the PPP attributes for dynamically created si interfaces:

1. Specify the predefined dynamic interface and logical interface variables in the dynamic profile.

```
[edit dynamic-profiles profile-name]  
user@host# edit interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
```

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set keepalives interval seconds
```

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set ppp-options chap  
user@host# set ppp-options pap
```

- Specify a set of AAA options that is used for authentication and authorization of tunneled PPP subscribers at the LNS that are logging in by means of the subscriber and AAA contexts that are specified in the AAA options set.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set ppp-options aaa-options aaa-options-name
```

The option set is configured with the **aaa-options** *aaa-options-name* statement at the **[edit access]** hierarchy level.

- Configure the router to prompt Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for tunneled PPP subscribers at the LNS.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set ppp-options ipcp-suggest-dns-option
```

- (Optional) Disable validation of the PPP magic number during LCP negotiation and in LCP keepalive (echo-request/echo-reply) exchanges. Prevents comparison of received magic number with internally generated magic number, so that a mismatch does not cause session termination.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set ppp-options ignore-magic-number-mismatch
```

To configure the PPP attributes for statically created si interfaces:

- Specify the logical inline service interface.

```
[edit interfaces si-slot/pic/port]
user@host# edit unit logical-unit-number
```

- Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

```
[edit interfaces si-slot/pic/port unit logical-unit-number]
user@host# set keepalives interval seconds
```

- Configure the number of keepalive packets a destination must fail to receive before the network takes down a link.

```
[edit interfaces si-slot/pic/port unit logical-unit-number]
user@host# set keepalives down-count number
```



NOTE: The **keepalives up-count** option is typically not used for subscriber management.

4. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

```
[edit interfaces si-slot/pic/port unit logical-unit-number]
user@host# set ppp-options chap
user@host# set ppp-options pap
```

5. Configure the router to prompt the Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for tunneled PPP subscribers at the LNS.

```
[edit interfaces si-slot/pic/port unit logical-unit-number]
user@host# set ppp-options ipcp-suggest-dns-option
```



BEST PRACTICE: Although all other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—are supported, they are typically not used for subscriber management. We recommend that you leave these other statements at their default values.



NOTE: You can also configure PPP attributes with a user group profile that applies the attributes to all subscribers with that profile on a LAC client. See [“Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile” on page 245](#) for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.



NOTE: When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring Subscriber Session Timeout Options](#)
- [Configuring Username Modification for Subscriber Sessions on page 179](#)
- [Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 136](#)

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

You can configure a user group profile that enables the LNS to apply PPP attributes to the PPP subscribers tunneled from the LAC. The user group profile is associated with clients (LACs) in the L2TP access profile. Consequently all subscribers handled by a given client share the same PPP attributes.

To configure a user group profile:

1. Create the profile.

```
[edit access]
user@host# edit group-profile profile-name
```

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

```
[edit access group-profile profile-name]
user@host# set ppp keepalive seconds
```



NOTE: Changes to the keepalive interval in a user group profile affect only new L2TP sessions that come up after the change. Existing sessions are not affected.

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

```
[edit access group-profile profile-name]
user@host# set ppp ppp-options chap
user@host# set ppp ppp-options pap
```

4. Specify a set of AAA options that is used for authentication and authorization of tunneled PPP subscribers at the LNS that are logging in by means of the subscriber and AAA contexts that are specified in the AAA options set.

```
[edit access group-profile profile-name]
user@host# set ppp ppp-options aaa-options aaa-options-name
```

The option set is configured with the `aaa-options aaa-options-name` statement at the `[edit access]` hierarchy level.

5. Configure the router to prompt the Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for tunneled PPP subscribers at the LNS.

```
[edit access group-profile profile-name]
user@host# set ppp ppp-options ipcp-suggest-dns-option
```

6. (Optional) Disable the Packet Forwarding Engine from performing a validation check for PPP magic numbers received from a remote peer in LCP keepalive

(Echo-Request/Echo-Reply) exchanges. This prevents PPP from terminating the session when the number does not match the value agreed upon during LCP negotiation. This capability is useful when the remote PPP peers include arbitrary magic numbers in the keepalive packets. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

```
[edit access group-profile profile-name]  
user@host# set ppp ppp-options ignore-magic-number-mismatch
```

7. Configure how long the PPP subscriber session can be idle before it is considered to have timed out.

```
[edit access group-profile profile-name]  
user@host# set ppp idle-timeout 200
```



NOTE: You can also configure PPP attributes on a per-interface basis. See [“Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface” on page 242](#) for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.



NOTE: When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

**Related
Documentation**

- [Configuring an L2TP Access Profile on the LNS on page 246](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- *Understanding Session Options for Subscriber Access*
- *Configuring Subscriber Session Timeout Options*
- [Configuring Username Modification for Subscriber Sessions on page 179](#)
- [Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 136](#)

Configuring an L2TP Access Profile on the LNS

Access profiles define how to validate Layer 2 Tunneling Protocol (L2TP) connections and session requests. Within each L2TP access profile, you configure one or more clients (LACs). The client characteristics are used to authenticate LACs with matching passwords,

and to establish attributes of the client tunnel and session. You can configure multiple access profiles and multiple clients within each profile.

To configure an L2TP access profile:

1. Create the access profile.

```
[edit access]
user@host# edit profile access-profile-name
```

2. Configure characteristics for one or more clients (LACs).

```
[edit access profile access-profile-name]
user@host# client client-name
```



NOTE: Except for the special case of the default client, the LAC client name that you configure in the access profile must match the hostname of the LAC. In the case of a Juniper Networks router acting as the LAC, the hostname is configured in the LAC tunnel profile with the gateway `gateway-name` statement at the `[edit access tunnel-profile profile-name tunnel tunnel-id source-gateway]` hierarchy level. Alternatively, the client name can be returned from RADIUS in the attribute, Tunnel-Client-Auth-Id [90].



NOTE: Use `default` as the client name when you want to define a default tunnel client. The default client enables the authentication of multiple LACs with the same secret and L2TP attributes. This behavior is useful when, for example, many new LACs are added to the network, because it enables the LACs to be used without additional LNS profile configuration.

Use `default` only on MX Series routers. The equivalent client name on M Series routers is `*`.

3. (Optional) Specify a local access profile that overrides the global access profile and the tunnel group AAA access profile to configure RADIUS server settings for the client.

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp aaa-access-profile
```

4. Configure the LNS to renegotiate the link control protocol (LCP) with the PPP client tunneled from the client.

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp lcp-renegotiation
```

5. Configure one or more dynamic service profiles to apply services to all subscribers on the LAC. You can optionally pass parameter to the services in the same statement.

```
[edit access profile access-profile-name client client-name]
```

```
user@host# set l2tp service-profile profile-name(parameter)&profile-name
```

6. Configure the maximum number of sessions allowed in a tunnel from the client (LAC).

```
[edit access profile access-profile-name client client-name]  
user@host# set l2tp maximum-sessions-per-tunnel number
```

7. Configure the LNS to override result codes 4 and 5 with result code 2 in CDN messages it sends to the LAC when the number of L2TP sessions reaches the configured maximum value. Some third-party LACs cannot fail over to another LNS unless the result code has a value of 2.

```
[edit access profile access-profile-name client client-name]  
user@host# set l2tp override-result-code session-out-of-resource
```

8. Configure the tunnel password used to authenticate the client (LAC).

```
[edit access profile access-profile-name client client-name]  
user@host# set l2tp shared-secret shared-secret
```

9. (Optional) Associate a group profile containing PPP attributes to apply for the PPP sessions being tunneled from this LAC client.

```
[edit access profile access-profile-name client client-name]  
user@host# set user-group-profile group-profile-name
```



NOTE: If *user-group-profile* is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.

**Related
Documentation**

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)

Configuring a AAA Local Access Profile on the LNS

For some LNS tunnels, you might wish to override the access profile configured at the routing instance that hosts the tunnel with a particular RADIUS server configuration. You can configure a local access profile to do so. You can subsequently use the **aaa-access-profile** statement to apply the local access profile to a tunnel group or LAC client.

A local access profile applied to a client overrides a local access profile applied to a tunnel group, which in turn overrides the access profile for the routing instance.

To configure an AAA local access profile:

1. Create the access profile.

```
[edit access]
user@host# edit profile local-aaa-profile-name
```

2. Configure the order of AAA authentication methods.

```
[edit access profile local-aaa-profile-name]
user@host# set authentication-order radius
```

3. Configure the RADIUS server attributes, such as the authentication password.

```
[edit access profile local-aaa-profile-name]
user@host# set radius-server server-address secret password
```

Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)

Configuring an Address-Assignment Pool for L2TP LNS with Inline Services

You can configure pools of addresses that can be dynamically assigned to the tunneled PPP subscribers. The pools must be local to the routing instance where the subscriber comes up. The configured pools are supplied in the RADIUS Framed-Pool and Framed-IPv6-Pool attributes. Pools are optional when Framed-IP-Address is sent by RADIUS.

To configure an address-assignment pool, you must specify the name of the pool and configure the addresses for the pool.

You can optionally configure multiple named ranges, or subsets, of addresses within an address-assignment pool. During dynamic address assignment, a client can be assigned an address from a specific named range. To create a named range, you specify a name for the range and define the address range.



NOTE: Be sure to use the address-assignment pools (**address-assignment**) statement rather than the address pools (**address-pool**) statement.

To configure an IPv4 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv4 family.

```
[edit access]
user@host# edit address-assignment pool pool-name family inet
```

2. Configure the network address and the prefix length of the addresses in the pool.

```
[edit access address-assignment pool pool-name family inet]
user@host# set network ip-prefix </prefix-length>
```

3. Configure the name of the range and the lower and upper boundaries of the addresses in the range.

```
[edit access address-assignment pool pool-name family inet]
user@host# set range range-name low lower-limit high upper-limit
```

For example, to configure an IPv4 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v4-pool family inet
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set network 192.168.1.1/16
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
```

To configure an IPv6 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv6 family.

```
[edit access]
user@host# edit address-assignment pool pool-name family inet6
```

2. Configure the IPv6 network prefix for the address pool. The prefix specification is required when you configure an IPv6 address-assignment pool.

```
[edit access address-assignment pool pool-name family inet6]
user@host# set prefix ipv6-prefix
```

3. Configure the name of the range and define the range. You can define the range based on the lower and upper boundaries of the prefixes in the range, or based on the length of the prefixes in the range.

```
[edit access address-assignment pool pool-name family inet6]
user@host# set range range-name low lower-limit high upper-limit
```

For example, to configure an IPv6 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v6-pool family inet6
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set prefix 2001:DB8::/32
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set range lns-v6-pool-range low 2001:DB8:1::/48 high 2001:DB8::ffff::/48
```

Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Address-Assignment Pools Overview](#)
- [Configuring Address-Assignment Pools](#)

Configuring the L2TP LNS Peer Interface

The peer interface connects the LNS to the cloud towards the LACs so that IP packets can be exchanged between the tunnel endpoints. MPLS and aggregated Ethernet can also be used to reach the LACs.



NOTE: On MX Series routers, you must configure the peer interface on an MPC.

To configure the LNS peer interface:

1. Specify the interface name.

```
[edit interfaces]
user@host# edit interface-name
```

2. Enable VLANs.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

3. Specify the logical interface, bind a VLAN tag ID to the interface, and configure the address family and the IP address for the logical interface.

```
[edit interfaces interface-name]
user@host# edit unit logical-unit-number
[edit interfaces interface-name unit logical-unit-number]
user@host# set vlan-id number
user@host# set family family address ip-address
```



NOTE: The IPv6 address family is not supported as a tunnel endpoint.

Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Enabling Inline Service Interfaces

The inline service interface is a virtual physical interface that resides on the Packet Forwarding Engine. This si interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.



NOTE: On MX80 and MX104 routers, you can configure only four inline services physical interfaces as anchor interfaces for L2TP LNS sessions: si-1/0/0, si-1/1/0, si-1/2/0, and si-1/3/0. You cannot configure si-0/0/0 for this purpose on MX80 and MX104 routers.

To enable inline service interfaces:

1. Access an MPC-occupied slot and the PIC where the interface is to be enabled.

```
[edit chassis]
user@host# edit fpc slot-number pic number
```

2. Enable the interface and specify the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.

```
[edit chassis fpc slot-number pic number]
user@host# set inline-services bandwidth (1g | 10g | 20g | 30g | 40g)
```

**Related
Documentation**

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Configuring an Inline Service Interface for L2TP LNS

The inline service interface is a virtual physical service interface that resides on the Packet Forwarding Engine. This *si* interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

You can maximize the number of sessions that can be shaped in one service interface by setting the maximum number of hierarchy levels to two. In this case, each LNS session consumes one L3 node in the scheduler hierarchy for shaping.

If you do not specify the number of levels (two is the only option), then the number of LNS sessions that can be shaped on the service interface is limited to the number of L2 nodes, or 4096 sessions. Additional sessions still come up, but they are not shaped.

To configure an inline service interface:

1. Access the service interface.

```
[edit interfaces]
user@host# edit si-slot/pic/port
```

2. (Optional; for per-session shaping only) Enable the inline service interface for hierarchical schedulers and limit the number of scheduler levels to two.

```
[edit interfaces si-slot/pic/port]
user@host# set hierarchical-scheduler maximum-hierarchy-levels 2
```


3. (Optional; for per-session shaping only) Configure services encapsulation for inline service interface.

```
[edit interfaces si-slot/pic/port]
user@host# set encapsulation generic-services
```

4. Configure the IPv4 family on the reserved unit 0 logical interface.

```
[edit interfaces si-slot/pic/port]
user@host# set unit 0 family inet
```

**Related
Documentation**

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Configuring Options for the LNS Inline Services Logical Interface

You must specify characteristics—**dial-options**—for each of the inline services logical interfaces that you configure for the LNS. LNS on MX Series routers supports only one session per logical interface, so you must configure it as a **dedicated** interface; the **shared** option is not supported. (LNS on M Series routers supports **dedicated** and **shared** options.) You also configure an identifying name for the logical interface that matches the name you specify in the access profile.

You must specify the **inet** address family for each static logical interface or in the dynamic profile for dynamic LNS interfaces. Although the CLI accepts either **inet** or **inet6** for static logical interfaces, the subscriber cannot log in successfully unless the address family **inet** is configured.



NOTE: For dynamic interface configuration, see “[Configuring a Dynamic Profile for Dynamic LNS Sessions](#)” on page 285.

To configure the static logical interface options:

1. Access the inline services logical interface.

```
[edit]
user@host# edit interfaces si-fpc/pic/port unit logical-unit-number
```

2. Specify an identifier for the logical interface.

```
[edit interfaces si-fpc/pic/port unit logical-unit-number]
user@host# set dial-options l2tp-interface-id name
```

3. Configure the logical interface to be used for only one session at a time.

```
[edit interfaces si-fpc/pic/port unit logical-unit-number]
user@host# set dial-options dedicated
```

4. Configure the address family for each logical interface and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.

```
[edit interfaces si-fpc/pic/port unit logical-unit-number]  
user@host# set family inet unnumbered-address lo0.0
```

**Related
Documentation**

- [Configuring a Dynamic Profile for Dynamic LNS Sessions on page 285](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)

LNS 1:1 Stateful Redundancy Overview

By default, when an inline service (si) anchor interface goes down—for example, when the card hosting the interface fails or restarts—L2TP subscriber traffic is lost. When the PPP keepalive timer for the tunnel subsequently expires, the control plane goes down and the PPP client is disconnected. Consequently, the client must then reconnect.

You can avoid traffic loss in these circumstances by configuring an aggregated inline service interface (asi) bundle to provide 1:1 stateful redundancy, also called hot standby or active-backup redundancy. The bundle consists of a pair of si physical interfaces, the primary (active) member link and the secondary (standby or backup) member link. These interfaces must be configured on different MPCs; redundancy is not achievable if you configure the primary and secondary interface on the same MPC because both member interfaces go down if the card goes down.

When subscribers log in and 1:1 redundancy is configured, the L2TP session is established over an underlying virtual logical interface (asix.0) over the asi0 physical interface. Individual subscriber logical interfaces are created on the underlying interface in the format, asiX.*logical-unit-number*. The session remains up in the event of a failure or a restart on the MPC hosting the primary member link interface. All the data traffic destined for this L2TP session automatically moves over to the secondary member link interface on the other MPC.

**Related
Documentation**

- [Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces

You can create an aggregated inline service interface (asi) bundle to provide 1:1 LNS stateful redundancy for inline service (si) anchor interfaces. The bundle pairs two interfaces that reside on different MPCs as primary and secondary links. LNS sessions are subsequently established over a virtual logical interface, `asiX.logical-unit-number`. LNS session failover occurs when either the primary anchor interface goes down or the card is restarted with the **request chassis fpc restart** command. When this happens, the secondary link—on a different MPC—becomes active and all the LNS data traffic destined for the session automatically moves over to the secondary interface. The subscriber session remains up on the `asiX.logical-unit-number` virtual interface. No traffic statistics are lost. When this redundancy is not configured, subscriber traffic is lost, the keepalives expire, and the PPP client is disconnected and must reconnect.

Before you begin, you must do the following:

- Confirm that enhanced subscriber management is enabled.
- Create inline service interfaces on different MPCs to be aggregated in the bundle.

See [“Enabling Inline Service Interfaces” on page 251](#) and [“Configuring an Inline Service Interface for L2TP LNS” on page 252](#).

- If you are using pools of service interfaces, define the service pools.



BEST PRACTICE: Follow these guidelines:

- You must configure unit 0 family inet for each bundle; otherwise, the session fails to come up.
- The primary (active) and secondary (backup) interfaces must be on different MPCs.
- The bandwidth configured at the `[edit chassis fpc slot pic number inline-services bandwidth]` hierarchy level must be the same for both member links.
- An si interface configured as a member of an aggregated inline service interface bundle cannot be configured as a member of another bundle group.
- An si interface configured as a member of an aggregated inline service interface bundle cannot also be used for any function that is not related to aggregated services; for example, it cannot be used for inline IP reassembly.
- When you configure an si interface as a member of an aggregated inline services bundle, you can no longer configure that si interface independently. You can configure only the parent bundle; the bundle's configuration is applied immediately to all member interfaces.

To configure 1:1 LNS stateful redundancy:

1. On one MPC, specify the primary (active) inline services member link in the bundle.

```
[edit interfaces asix aggregated-inline-services-options]
user@host# set primary-interface
```

2. Configure the amount of bandwidth reserved on this MPC for tunnel traffic using the primary inline service interface.

```
[edit chassis fpc slot pic number inline-services]
user@host# set bandwidth (1g | 10g)
```

3. On a different MPC, specify the secondary (backup) inline services member link in the bundle.

```
[edit interfaces asix aggregated-inline-services-options]
user@host# set secondary-interface
```



NOTE: If you configure the active and backup member links on the same MPC, the subsequent commit of the configuration fails.

4. Configure the amount of bandwidth reserved on this MPC for tunnel traffic using the secondary inline service interface.

```
[edit chassis fpc slot pic number inline-services]
user@host# set bandwidth (1g | 10g)
```

5. Assign the aggregated inline service interface bundle to an L2TP tunnel group by either of the following methods:

- Assign a single bundle by specifying the name of the aggregated inline service physical interface.

```
[edit services l2tp tunnel-group name]
user@host# set service-interface interface-name
```

- Assign one or more pools of bundles to the tunnel group.

```
[edit services l2tp tunnel-group name]
user@host# set service-device-pool pool-name
```



NOTE: A pool can be mixed; that is, it can include both aggregated inline service interface bundles and individual inline service interfaces. The individual interfaces must not be members of existing bundles.

The following sample configuration creates bundle asi0 with member links on MPCs in slot 1 and slot 2, then assigns the bundle to provide redundancy for L2TP sessions on tunnel group tg1:

```
[edit interfaces asi0]
```

```

user@host# set aggregated-inline-services-options primary-interface si-1/0/0
user@host# set aggregated-inline-services-options secondary-interface si-2/0/0
user@host# set unit 0 family inet

```

```

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

```

```

[edit chassis fpc 2 pic 0 inline-services]
user@host# set bandwidth 10g

```

```

[edit services l2tp tunnel-group tg1]
user@host# set service-interface asi0

```

Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

Verifying LNS Aggregated Inline Service Interface 1:1 Redundancy

Purpose View information about aggregated inline service interface bundles, individual member links, and redundancy status.

Action

- To view summary information about an aggregated inline service interface bundle:

```

user@host> show interfaces asi0 terse

```

Interface	Admin	Link	Proto	Local	Remote
asi0	up	up			
asi0.0	up	up	inet		

- To view detailed information about an aggregated inline service interface bundle:

```

user@host> show interfaces asi0 extensive
Physical interface: asi0, Enabled, Physical link is Up
  Interface index: 223, SNMP ifIndex: 734, Generation: 226
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 20000mbps
  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   : 2014-01-20 23:35:02 PST (00:03:25 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets : 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets : 0
  Output packets: 0

```

```

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 asi0.0 (Index 356) (SNMP ifIndex 52241) (Generation 165)
Flags: Point-To-Point SNMP-Traps 0x4000 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
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 9192, Generation: 198, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

- To view information about an individual member interface in an aggregated inline service interface bundle:

```

user@host> show interfaces si-1/0/0
Physical interface: si-1/0/0, Enabled, Physical link is Up
Interface index: 165, SNMP ifIndex: 630
Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192, Speed:
10000mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Link type : Full-Duplex
Link flags : None
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface si-1/0/0.0 (Index 357) (SNMP ifIndex 52229)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Adaptive-Services
Input packets : 0
Output packets: 0
Protocol asi, AS bundle: asi0.0
Flags: Function2

```

- To view redundancy status for aggregated inline service interface bundles:

```

user@host> show interfaces redundancy

```

Interface	State	Last change	Primary	Secondary	Current status
asi0	On secondary	1d 23:56	si-1/0/0	si-2/0/0	primary down
asi1	On primary	10:10:27	si-3/0/0	si-4/0/0	secondary down
ae0	On primary	00:00:02	ge-1/0/0	ge-3/0/1	backup down
ae2	On primary	00:00:01	ge-2/0/0	ge-4/0/1	both up

That sample output shows that both aggregated Ethernet and aggregated inline service interfaces are configured for redundancy. To display only one of the aggregated inline service interface bundles:

```
user@host> show interfaces redundancy asi0
Interface State      Last change Primary   Secondary Current status
asi0       On secondary 1d 23:56   si-1/0/0  si-2/0/0  primary down
```

- To view detailed information about all configured redundancy interfaces:

```
user@host> show interfaces redundancy detail
Redundancy interfaces detail
Interface      : asi0
State          : On primary
Last change    : 00:00:36
Primary        : si-1/0/0
Secondary      : si-3/0/0
Current status : both up

Interface      : ae0
State          : On primary
Last change    : 00:01:30
Primary        : ge-1/0/0
Secondary      : ge-3/0/1
Current status : backup down
```

Related Documentation

- [Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [LNS 1:1 Stateful Redundancy Overview on page 254](#)

L2TP Session Limits and Load Balancing for Service Interfaces

The LNS load balances subscriber sessions across the available service interfaces in a device pool based on the number of sessions currently active on the interfaces. You can configure a maximum limit per service interface (si) and per aggregated service interface (asi). In the case of asi interfaces, you cannot configure a limit for the individual si member interfaces in the bundle.

Session Limits on Service Interfaces

When an L2TP session request is initiated for a service interface, the LNS checks the number of current active sessions on that interface against the maximum number of sessions allowed for the individual service interface or aggregated service interface. The LNS determines whether the current session count (displayed by the **show services l2tp summary** command) is less than the configured limit. When that is true or when no limit is configured, the check passes and the session can be established. If the current session count is equal to the configured limit, then the LNS rejects the session request. No subsequent requests can be accepted on that interface until the number of active requests drops below the configured maximum. When a session request is rejected for an si or asi

interface, the LNS returns a CDN message with the result code set to 2 and the error code set to 4.

For example, suppose a single service interface is configured in the tunnel group. The current L2TP session count is 1500, with a configured limit of 2000 sessions. When a new session is requested, the limit check passes and the session request is accepted.

Interface	Configured Session Limit	Current Session Count	Session Limit Check Result
si-0/0/0	2000	1500	Pass

The limit check continues to pass and session requests are accepted until 500 requests have been accepted, making the current session count 2000, which matches the configured maximum. The session limit check fails for all subsequent requests and all requests are rejected until the current session count on the interface drops below 2000, so that the limit check can pass.

Interface	Configured Session Limit	Current Session Count	Session Limit Check Result
si-0/0/0	2000	2000	Fail

When the session limit is set to zero for an interface, no session requests can be accepted. If that is the only interface in the tunnel group, then all session requests in the group are rejected until the session limit is increased from zero or another service interface is added to the tunnel group.

When a service interface in a service device pool has reached the maximum configured limit or it has a configured limit of zero, the LNS skips that interface when a session request is made and selects another interface in the pool to check the session limit. This continues until an interface passes and the session is accepted or no other interface remains in the pool to be selected.

Session Load Balancing Across Service Interfaces

The behavior for session load distribution in a service device pool changed in Junos OS Release 16.2. When a service interface has a lower session count than another interface in the pool and both interfaces are below their maximum session limit, subsequent sessions are distributed to the interface with fewer sessions.

In earlier releases, sessions are distributed in a strictly round-robin manner, regardless of session count. The old behavior can result in uneven session distribution when the Packet Forwarding Engine is rebooted or a service interface goes down and comes back up.

For example, consider the following scenario using the old round-robin distribution behavior for a pool with two service interfaces:

1. Two hundred sessions are evenly distributed across the two service interfaces.
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 100 sessions.
2. The si-1/0/0 interface reboots. When it comes back, initially sessions are up only on si-0/0/0.
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 0 sessions.
3. As the sessions formerly on si-1/0/0 reconnect, they are distributed equally across both service interfaces. When all 100 sessions are back up, the distribution is significantly unbalanced.
 - si-0/0/0 has 150 sessions.
 - si-1/0/0 has 50 sessions.
4. After 100 new sessions connect, si-0/0/0 reaches its maximum limit. Subsequent sessions are accepted only on si-1/0/0.
 - si-0/0/0 has 200 sessions.
 - si-1/0/0 has 100 sessions.
5. After 100 more sessions connect, si-1/0/0 reaches its maximum limit. No more sessions can be accepted until the session count drops below 200 for one of the interfaces.
 - si-0/0/0 has 200 sessions.
 - si-1/0/0 has 200 sessions.

Now consider the same scenario using the current load distribution behavior based on the number of attached sessions. The device pool again has two service interfaces each with a configured maximum limit of 200 sessions:

1. Two hundred sessions are evenly distributed across the two service interfaces.
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 100 sessions.
2. The si-1/0/0 interface reboots. When it comes back up, sessions are up initially only on si-0/0/0.
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 0 sessions.
3. As the sessions formerly on si-1/0/0 reconnect, they are distributed according to the session load on each interface. Because both interfaces are below their maximum limit, and si-1/0/0 has fewer sessions than si-0/0/0, sessions are initially distributed only to si-1/0/0.
 - a. After 1 new session:

- si-0/0/0 has 100 sessions.
- si-1/0/0 has 1 session.
- b. After 10 new sessions:
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 10 sessions.
- c. After 100 new sessions:
 - si-0/0/0 has 100 sessions.
 - si-1/0/0 has 100 sessions.
- 4. Because both interfaces now have the same session count, the next session (#101) is distributed randomly between the two interfaces. The next session after that (#102) goes to the interface with the lower session count. That makes the interfaces equal again, so the next session (#103) is randomly distributed. This pattern repeats until the maximum limit of 200 sessions for both interfaces.
 - si-0/0/0 has 200 sessions.
 - si-1/0/0 has 200 sessions.

No more sessions can be accepted on either interface until the number of sessions drops below 200 on one of the interfaces.

The load balancing behavior is the same for aggregated service interfaces. An asi interface is selected from a pool based on the current session count for the asi interface. When that count is less than the maximum, the LNS checks current session count for the active si interface in the asi bundle. When that count is less than the maximum, the session can be established on the asi interface.

In a mixed device pool that has both service interfaces and aggregated service interfaces, sessions are distributed to the interface, either asi or si, that has the lowest session count. When the session count of an interface of either type reaches its limit, it can no longer accept sessions until the count drops below the maximum.

You can use the session limit configuration to achieve a session limit on particular Packet Forwarding Engines. Suppose you want a limit of 100 sessions on a PFE0, which has two service interfaces. You can set the max limit on each interface to 50, or any other combination that adds up to 100 to establish the PFE0 limit.

**Related
Documentation**

- [L2TP Session Limits Overview on page 199](#)
- [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204](#)
- [L2TP for Subscriber Access Overview on page 147](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP LAC on page 173](#)

Example: Configuring an L2TP LNS

This example shows how you can configure an L2TP LNS on an MX Series router to provide tunnel endpoints for an L2TP LAC in your network. This configuration includes a dynamic profile for dual-stack subscribers.

- [Requirements on page 263](#)
- [Overview on page 264](#)
- [Configuration on page 265](#)

Requirements

This L2TP LNS example requires the following hardware and software:

- MX Series 3D Universal Edge Router
- One or more MPCs
- Junos OS Release 11.4 or later

No special configuration beyond device initialization is required before you can configure this feature.

You must configure certain standard RADIUS attributes and Juniper Networks VSAs in the attribute return list on the AAA server associated with the LNS for this example to work. [Table 24 on page 263](#) lists the attributes with their required order setting and values. We recommend that you use the most current Juniper Networks RADIUS dictionary, available in the *Downloads* box on the *Junos OS Subscriber Management* page at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/index.html.

Table 24: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example

VSA Name [Number]	Order	Value
CoS-Parameter-Type [26-108]	1	T01 Multiplay
CoS-Parameter-Type [26-108]	2	T02 10m
CoS-Parameter-Type [26-108]	3	T08 -36
CoS-Parameter-Type [26-108]	4	T07 cell-mode
Framed-IPv6-Pool [100]	0	jnpr_ipv6_pool
Framed-Pool [88]	0	jnpr_pool
Egress-Policy-Name [26-11]	0	classify
Ingress-Policy-Name [26-10]	0	classify
Virtual-Router [26-1]	0	default

Overview

The LNS employs user group profiles to apply PPP attributes to the PPP subscribers that are tunneled from the LAC. LACs in the network are clients of the LNS. The clients are associated with user group profiles in the L2TP access profile configured on the LNS. In this example, the user group profile **ce-l2tp-group-profile** specifies the following PPP attributes:

- A 30-second interval between PPP keepalive messages for L2TP tunnels from the client LAC terminating on the LNS.
- A 200-second interval that defines how long the PPP subscriber session can be idle before it is considered to have timed out.
- Both PAP and CHAP as the PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

The L2TP access profile **ce-l2tp-profile** defines a set of L2TP parameters for each client LAC. In this example, the user group profile **ce-l2tp-group-profile** is associated with both clients, **lac1** and **lac2**. Both clients are configured to have the LNS renegotiate the link control protocol (LCP) with the PPP client rather than accepting the pre-negotiated LCP parameters that the LACs pass to the LNS. LCP renegotiation also causes authentication to be renegotiated by the LNS; the authentication method is specified in the user group profile. The maximum number of sessions allowed per tunnel is set to 1000 for **lac1** and to 4000 for **lac2**. A different password is configured for each LAC.

A local AAA access profile, **aaa-profile**, enables you to override the global AAA access profile, so that you can specify an authentication order, a RADIUS server that you want to use for L2TP, and a password for the server.

In this example, an address pool defines a range of IP addresses that the LNS allocates to the tunneled PPP sessions. This example defines ranges of IPv4 and IPv6 addresses.

Two inline service interfaces are enabled on the MPC located in slot 5 of the router. For each interface, 10 Gbps of bandwidth is reserved for tunnel traffic on the interface's associated PFE. These *anchor* interfaces serve as the underlying physical interface. To enable CoS queue support on the individual logical inline service interfaces, you must configure both services encapsulation (**generic-services**) and hierarchical scheduling support on the anchors. The IPv4 address family is configured for both anchor interfaces. Both anchor interfaces are specified in the **lns_p1** service device pool. The LNS can balance traffic loads across the two anchor interfaces when the tunnel group includes the pool.

This example uses the dynamic profile **dyn-ns-profile2** to specify characteristics of the L2TP sessions that are created or assigned dynamically when a subscriber is tunneled to the LNS. For many of the characteristics, a predefined variable is set; the variables are dynamically replaced with the appropriate values when a subscriber is tunneled to the LNS.

The interface to which the tunneled PPP client connects (**\$junos-interface-name**) is dynamically created in the routing instance (**\$junos-routing-instance**) assigned to the subscriber. Routing options for access routes include the route's next hop address

(**\$junos-framed-route-nexthop**), metric (**\$junos-framed-route-cost**), and preference (**\$junos-framed-route-distance**). For access-internal routes, a dynamic IP address variable (**\$junos-subscriber-ip-address**) is set.

The logical inline service interfaces are defined by the name of a configured anchor interface (**\$junos-interface-ifd-name**) and a logical unit number (**\$junos-interface-unit**). The profile assigns **l2tp-encapsulation** as the identifier for the logical interface and specifies that each interface can be used for only a single session at a time.

The IPv4 address is set to a value returned from the AAA server. For IPv4 traffic an input firewall filter **\$junos-input-filter** and an output firewall filter **\$junos-output-filter** are attached to the interface. The loopback variable (**\$junos-loopback-interface**) derives an IP address from a loopback interface (**lo**) configured in the routing instance and uses it in IPCP negotiation as the PPP server address. Because this is a dual-stack configuration, the IPv6 address family is also set, with the addresses provided by the **\$junos-ipv6-address** variable.

The **\$junos-ipv6-address** variable is used because Router Advertisement Protocol is also configured. This variable enables AAA to allocate the first address in the prefix to be reserved as the local address for the interface. The minimal configuration for the Router Advertisement Protocol in the dynamic profile specifies the **\$junos-interface-name** and **\$junos-ipv6-ndra-prefix** variables to dynamically assign a prefix value in IPv6 neighbor discovery router advertisements.

The dynamic profile also includes the class of service configuration that is applied to the tunnel traffic. The traffic-control profile (**tc-profile**) includes variables for the scheduler map (**\$junos-cos-scheduler-map**), shaping rate (**\$junos-cos-shaping-rate**), overhead accounting (**\$junos-cos-shaping-mode**), and byte adjustment **\$junos-cos-byte-adjust**. The dynamic profile applies the CoS configuration—including the forwarding class, the output traffic-control profile, and the rewrite rules—to the dynamic service interfaces.

The **tg-dynamic** tunnel group configuration specifies the access profile **ce-l2tp-profile**, the local AAA profile **aaa-profile**, and the dynamic profile **dyn-lns-profile2** that are used to dynamically create LNS sessions and define the characteristics of the sessions. The **lns_p1** service device pool associates a pool of service interfaces with the group to enable LNS to balance traffic across the interfaces. The local gateway address **203.0.113.2** corresponds to the remote gateway address that is configured on the LAC. The local gateway name **ce-lns** corresponds to the remote gateway name that is configured on the LAC.



NOTE: This example does not show all possible configuration choices.

Configuration

CLI Quick Configuration

To quickly configure an L2TP LNS, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
[edit]
edit access group-profile ce-l2tp-group-profile
```

```
set ppp idle-timeout 200
set ppp ppp-options pap
set ppp ppp-options chap
set ppp keepalive 30
top
edit access profile ce-l2tp-profile
set client lac1 l2tp maximum-sessions-per-tunnel 1000
set client lac1 l2tp interface-id l2tp-encapsulation-1
set client lac1 l2tp lcp-renegotiation
set client lac1 l2tp shared-secret "lac1-$ABC123"
set client lac1 user-group-profile ce-l2tp-group-profile
set client lac2 l2tp maximum-sessions-per-tunnel 4000
set client lac2 l2tp interface-id l2tp-encap-2
set client lac2 l2tp lcp-renegotiation
set client lac2 l2tp shared-secret "lac2-$ABC123"
set client lac2 user-group-profile ce-l2tp-group-profile
top
edit access profile aaa-profile
set authentication-order radius
set radius authentication-server 198.51.100.193
set radius-server 198.51.100.193 secret "$ABC123"
top
edit access address-assignment pool client-pool1 family inet
set network 192.168.1.1/16
set range lns-v4-pool-range low 192.168.1.1
set range lns-v4-pool-range high 192.168.255.255
top
edit access address-assignment pool client-ipv6-pool2 family inet6
set prefix 2001:DB8::/32
set range lns-v6-pool-range low 2001:DB8:1::/48
set range lns-v6-pool-range high 2001:DB8:ffff::/48
top
set interfaces ge-5/0/1 unit 11 vlan-id 11
set interfaces ge-5/0/1 unit 11 family inet address 203.0.113.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
top
set chassis fpc 5 pic 0 inline-services bandwidth 10g
set chassis fpc 5 pic 2 inline-services bandwidth 10g
top
edit interfaces si-5/0/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
edit interfaces si-5/2/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
set services service-device-pools pool lns_p1 interface si-5/0/0
set services service-device-pools pool lns_p1 interface si-5/2/0
top
edit dynamic-profiles dyn-lns-profile2 routing-instances $junos-routing-instance
set interface $junos-interface-name
edit routing-options access route $junos-framed-route-ip-address-prefix
set next-hop $junos-framed-route-nexthop
```

```

set metric $junos-framed-route-cost
set preference $junos-framed-route-distance
up 2
edit access-internal route $junos-subscriber-ip-address
set qualified-next-hop $junos-interface-name
up 5
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set dial-options l2tp-interface-id l2tp-encapsulation
set dial-options dedicated
set family inet filter input $junos-input-filter
set family inet filter output $junos-output-filter
set family inet unnumbered-address $junos-loopback-interface
set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
up 3
edit protocols router-advertisement
set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix
top
[edit class-of-service]
edit rewrite-rules dscp rewriteDSCP forwarding-class expedited-forwarding
set loss-priority high code-point af11
set loss-priority high code-point af12
top
edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles tc-profile
set scheduler-map $junos-cos-scheduler-map
set shaping-rate $junos-cos-shaping-rate
set overhead-accounting $junos-cos-shaping-mode
set overhead-accounting bytes $junos-cos-byte-adjust
up
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set forwarding-class expedited-forwarding
set output-traffic-control-profile tc-profile
set rewrite-rules dscp rewriteDSCP
edit interfaces si-5/0/0
set output-control-profile-remaining tc-profile
top
set services l2tp tunnel-group tg-dynamic l2tp-access-profile ce-l2tp-profile
set services l2tp tunnel-group tg-dynamic aaa-access-profile aaa-profile
set services l2tp tunnel-group tg-dynamic local-gateway address 203.0.113.2
set services l2tp tunnel-group tg-dynamic local-gateway gateway-name ce-lns
set services l2tp tunnel-group tg-dynamic service-device-pool lns_p1
set services l2tp tunnel-group tg-dynamic dynamic-profile dyn-lns-profile2

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an L2TP LNS with inline service interfaces:

1. Configure a user group profile that defines the PPP configuration for tunnel subscribers.

[edit access]

```

user@host# edit group-profile ce-l2tp-group-profile
[edit access group-profile ce-l2tp-group-profile]
user@host# set ppp keepalive 30
user@host# set ppp idle-timeout 200
user@host# set ppp ppp-options chap
user@host# set ppp ppp-options pap

```

2. Configure an L2TP access profile that defines the L2TP parameters for each client LAC. This includes associating a user group profile with the client and specifying the identifier for the inline services logical interface that represents an L2TP session on the LNS.

```

[edit access profile ce-l2tp-profile client lac1]
user@host# set l2tp interface-id l2tp-encapsulation
user@host# set l2tp maximum-sessions-per-tunnel 1000
user@host# set l2tp shared-secret "lac1-$ABC123"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile
[edit access profile ce-l2tp-profile client lac2]
user@host# set l2tp interface-id interface-id
user@host# set l2tp maximum-sessions-per-tunnel 4000
user@host# set l2tp shared-secret "lac2-$ABC123"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile

```



NOTE: If `user-group-profile` is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.

3. Configure a AAA access profile to override the global access profile for the order of AAA authentication methods and server attributes.

```

[edit access profile aaa-profile]
user@host# set authentication-order radius
user@host# set radius authentication-server 198.51.100.193
user@host# set radius-server 198.51.100.193 secret "$ABC123"

```

4. Configure IPv4 and IPv6 address-assignment pools to allocate addresses for the clients (LACs).

```

[edit access address-assignment pool client-pool1 family inet]
user@host# set network 192.168.1.1/16
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
[edit access address-assignment pool client-ipv6-pool2 family inet6]
user@host# set prefix 2001:DB8::/32
user@host# set range lns-v6-pool-range low 2001:DB8:1::/48
user@host# set range lns-v6-pool-range high 2001:DB8:ffff::/48

```

5. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address (loopback address).


```
[edit interfaces ge-5/0/1
user@host# set vlan-tagging
user@host# set unit 11
[edit interfaces ge-5/0/1.11
user@host# set vlan-id 11
user@host# set family inet address 10.1.1.2/24
[edit interfaces lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
```

6. Enable inline service interfaces on an MPC.

```
[edit chassis fpc 5]
user@host# set pic 0 inline-services bandwidth 10g
user@host# set pic 2 inline-services bandwidth 10g
```

7. Configure the anchor service interfaces with services encapsulation, hierarchical scheduling, and the address family.

```
[edit interfaces si-5/0/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
[edit interfaces si-5/2/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
```

8. Configure a pool of service interfaces for dynamic LNS sessions.

```
[edit services service-device-pools pool lns_p1]
user@host# set interface si-5/0/0
user@host# set interface si-5/2/0
```

9. Configure a dynamic profile that dynamically creates L2TP logical interfaces for dual-stack subscribers.

```
[edit dynamic-profiles dyn-lns-profile2]
user@host# edit routing-instances $junos-routing-instance
user@host# set interface $junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"]
user@host# edit routing-options access route $junos-framed-route-ip-address-prefix
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
routing-options access route "$junos-framed-route-ip-address-prefix"]
user@host# set next-hop $junos-framed-route-nexthop
user@host# set metric $junos-framed-route-cost
user@host# set preference $junos-framed-route-distance
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
routing-options access-internal]
user@host# set route $junos-subscriber-ip-address qualified-next-hop
$junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set dial-options l2tp-interface-id l2tp-encapsulation
user@host# set dial-options dedicated
```

```

user@host# set family inet unnumbered-address $junos-loopback-interface
user@host# set family inet filter input $junos-input-filter
user@host# set family inet filter output $junos-output-filter
user@host# set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
[edit dynamic-profiles dyn-lns-profile2 protocols router-advertisement]
user@host# set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix

```

10. Configure shaping, scheduling, and rewrite rules, and apply in the dynamic profile to tunnel traffic.

```

[edit class-of-service]
user@host# edit rewrite-rules dscp rewriteDSCP forwarding-class
expedited-forwarding
user@host# set loss-priority high code-point af11
user@host# set loss-priority high code-point af12
[edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles
tc-profile]
user@host# set scheduler-map $junos-cos-scheduler-map
user@host# set shaping-rate $junos-cos-shaping-rate
user@host# set overhead-accounting $junos-cos-shaping-mode
user@host# set overhead-accounting bytes $junos-cos-byte-adjust
[edit dynamic-profiles dyn-lns-profile2 class-of-service interfaces
"$junos-interface-ifd-name" unit "$junos-interface-unit"]
user@host# set forwarding-class expedited-forwarding
user@host# set output-traffic-control-profile tc-profile
user@host# set rewrite-rules dscp rewriteDSCP
[edit class-of-service interfaces si-5/0/0]
user@host# set output-traffic-control-profile-remaining tc-profile

```

11. Configure the L2TP tunnel group to bring up dynamic LNS sessions using the pool of inline service interfaces to enable load-balancing.

```

[edit services l2tp tunnel-group tg-dynamic]
user@host# set l2tp-access-profile ce-l2tp-profile
user@host# set local-gateway address 10.1.1.2
user@host# set local-gateway gateway-name ce-lns
user@host# set aaa-access-profile aaa-profile
user@host# set dynamic-profile dyn-lns-profile2
user@host# set service-device-pool lns_p1

```

Results From configuration mode, confirm the access profile, group profile, AAA profile, and address-assignment pools configuration by entering the **show access** command. Confirm the inline services configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. Confirm the dynamic profile configuration by entering the **show dynamic-profiles** command. Confirm the tunnel group configuration by entering the **show services l2tp** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

[edit]

```

user@host# show access
group-profile ce-l2tp-group-profile {
  ppp {
    idle-timeout 200;
    ppp-options {
      pap;
      chap;
    }
    keepalive 30;
  }
}
profile ce-l2tp-profile {
  client lac1 {
    l2tp {
      maximum-sessions-per-tunnel 1000;
      interface-id l2tp-encapsulation-1;
      lcp-renegotiation;
      shared-secret "lac1-$ABC123"; ## SECRET-DATA
    }
    user-group-profile ce-l2tp-group-profile;
  }
  client lac2 {
    l2tp {
      maximum-sessions-per-tunnel 4000;
      interface-id l2tp-encap-2;
      lcp-renegotiation;
      shared-secret "lac2-$ABC123"; ## SECRET-DATA
    }
    user-group-profile ce-l2tp-group-profile;
  }
}
profile aaa-profile {
  authentication-order radius;
  radius-server {
    198.51.100.193 secret "$ABC123"; ## SECRET-DATA
  }
}
address-assignment {
  pool client-pool1 {
    family inet {
      network 192.168.1.1/16;
      range lns-v4-pool-range {
        low 192.168.1.1;
        high 192.168.255.255;
      }
    }
  }
  pool client-ipv6-pool2 {
    family inet6 {
      prefix 2001:DB8::/32;
      range lns-v6-pool-range {
        low 2001:DB8:1::/48;
        high 2001:DB8:ffff::/48;
      }
    }
  }
}

```

```
}
```

```
[edit]
```

```
user@host# show chassis
```

```
fpc 5 {  
  pic 0 {  
    inline-services {  
      bandwidth 10g;  
    }  
  }  
  pic 2 {  
    inline-services {  
      bandwidth 10g;  
    }  
  }  
}
```

```
[edit]
```

```
user@host# show interfaces
```

```
ge-5/0/1 {  
  vlan-tagging;;  
  unit 11 {  
    vlan-id 11;  
    family inet {  
      address 203.0.113.2/24;  
    }  
  }  
}  
si-5/0/0 {  
  hierarchical-scheduler maximum-hierarchy-levels 2;  
  encapsulation generic-services;  
  unit 0 {  
    family inet;  
  }  
}  
si-5/2/0 {  
  hierarchical-scheduler maximum-hierarchy-levels 2;  
  encapsulation generic-services;  
  unit 0 {  
    family inet;  
  }  
}  
lo0 {  
  unit 0 {  
    family inet {  
      address 127.0.0.1/32;  
    }  
  }  
}
```

```
[edit]
```

```
user@host# show dynamic-profiles
```

```
dyn-lns-profile2 {  
  routing-instances {  
    "$junos-routing-instance" {  
      interface "$junos-interface-name";  
    }  
  }  
}
```

```

routing-options {
  access {
    route $junos-framed-route-ip-address-prefix {
      next-hop "$junos-framed-route-nexthop";
      metric "$junos-framed-route-cost";
      preference "$junos-framed-route-distance";
    }
  }
  access-internal {
    route $junos-subscriber-ip-address {
      qualified-next-hop "$junos-interface-name";
    }
  }
}
}
interfaces {
  "$junos-interface-ifd-name" {
    unit "$junos-interface-unit" {
      dial-options {
        l2tp-interface-id l2tp-encapsulation;
        dedicated;
      }
      family inet {
        filter {
          input "$junos-input-filter";
          output "$junos-output-filter";
        }
        unnumbered-address "$junos-loopback-interface";
      }
      family inet6 {
        address $junos-ipv6-address;
        input $junos-input-ipv6-filter;
        output $junos-output-ipv6-filter;
      }
    }
  }
}
protocols {
  router-advertisement {
    interface "$junos-interface-name" {
      prefix $junos-ipv6-ndra-prefix;
    }
  }
}
}
class-of-service {
  rewrite-rules {
    dscp rewriteDSCP {
      forwarding-class expedited-forwarding {
        loss-priority high code-point af11
        loss-priority high code-point af12
      }
    }
  }
}
}
traffic-control-profiles {
  tc-profile {

```

```
        scheduler-map "$junos-cos-scheduler-map";
        shaping-rate "$junos-cos-shaping-rate";
        overhead-accounting "$junos-cos-shaping-mode" bytes "$junos-cos-byte-adjust";
    }
}
interfaces {
    "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
            forwarding-class expedited-forwarding;
            output-traffic-control-profile tc-profile;
            rewrite-rules {
                dscp rewriteDSCP;
            }
        }
    }
}
}
}

[edit]
user@host# show services l2tp
tunnel-group tg-dynamic {
    l2tp-access-profile ce-l2tp-profile;
    aaa-access-profile aaa-profile;
    local-gateway {
        address 203.0.113.2;
        gateway-name ce-lns;
    }
    service-device-pool lns_p1;
    dynamic-profile dyn-lns-profile2;
}
```

When you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [L2TP for Subscriber Access Overview on page 147](#)
 - [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
 - [Configuring an L2TP LAC on page 173](#)

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

The L2TP tunnel group specifies attributes that apply to L2TP tunnels and sessions from a group of LAC clients. These attributes include the access profile used to validate L2TP connection requests made to the LNS on the local gateway address, a local access profile that overrides the global access profile, the keepalive timer, and whether the IP ToS value is reflected.



NOTE: If you delete a tunnel group, all L2TP sessions in that tunnel group are terminated. If you change the value of the `local-gateway-address`, `service-device-pool`, or `service-interface` statements, all L2TP sessions using those settings are terminated. If you change or delete other statements at the `[edit services l2tp tunnel-group name]` hierarchy level, new tunnels you establish use the updated values but existing tunnels and sessions are not affected.

To configure the LNS tunnel group:

1. Create the tunnel group.

```
[edit services l2tp]
user@host# edit tunnel-group group-name
```



NOTE: You can create up to 256 tunnel groups.

2. Specify the service anchor interface responsible for L2TP processing on the LNS.

```
[edit services l2tp tunnel-group name]
user@host# set service-interface interface-name
```

This service anchor interface is required for static LNS sessions, and for dynamic LNS sessions that do not balance traffic across a pool of anchor interfaces. The interface is configured at the `[edit interfaces]` hierarchy level.

3. (Optional; for load-balancing dynamic LNS sessions only) Specify a pool of inline service anchor interfaces to enable load-balancing of L2TP traffic across the interfaces.

```
[edit services l2tp tunnel-group group-name]
user@host# set service-device-pool pool-name
```

The pool is defined at the `[edit services service-device-pools]` hierarchy level.

4. (For dynamic LNS sessions only) Specify the name of the dynamic profile that defines and instantiates inline service interfaces for L2TP tunnels

```
[edit services l2tp tunnel-group group-name]
user@host# set dynamic-profile profile-name
```

The profile is defined at the `[edit dynamic-profiles]` hierarchy level.

5. Specify the access profile that validates all L2TP connection requests to the local gateway address.

```
[edit services l2tp tunnel-group group-name]
user@host# set l2tp-access-profile profile-name
```

6. Configure the local gateway address on the LNS; corresponds to the IP address that is used by LACs to identify the LNS.

```
[edit services l2tp tunnel-group group-name]  
user@host# set local-gateway address address
```

7. (Optional) Configure the local gateway name on the LNS, returned in the SCCRP message to the LAC. The name must match the remote gateway name configured on the LAC, or the tunnel cannot be created.

```
[edit services l2tp tunnel-group group-name]  
user@host# set local-gateway gateway-name gateway-name
```

8. (Optional) Configure the interval at which the LNS sends hello messages if it has received no messages from the LAC.

```
[edit services l2tp tunnel-group group-name]  
user@host# set hello-interval seconds
```

9. (Optional) Specify a local access profile that overrides the global access profile to configure RADIUS server settings for the tunnel group.

```
[edit services l2tp tunnel-group group-name]  
user@host# set aaa-access-profile profile-name
```

This local profile is configured at the **[edit access profile]** hierarchy level.

10. (Optional) Configure the LNS to reflect the IP ToS value from the inner IP header to the outer IP header (applies to CoS configurations).

```
[edit services l2tp tunnel-group group-name]  
user@host# set tos-reflect
```

11. (Optional) Configure a service profile to be applied to the L2TP session at login, along with any parameters to pass to the service.

```
[edit services l2tp tunnel-group group-name]  
user@host# set service-profile profile-name(parameter)&profile-name
```

**Related
Documentation**

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)
- [Applying Services to an L2TP Session Without Using RADIUS on page 277](#)

Applying Services to an L2TP Session Without Using RADIUS

Services are applied to L2TP sessions for activation or later modified by vendor-specific attributes (VSAs) from the RADIUS server or in RADIUS Change of Authorization (CoA) requests. Starting in Junos OS Release 18.1R1, you can apply services to L2TP sessions by means of dynamic service profiles without involving RADIUS. In multivendor environments, customers might use only standard RADIUS attributes to simplify management by avoiding the use of VSAs from multiple vendors. However, this complicates the application of services to L2TP sessions because VSAs are generally required to apply services. Local dynamic service profile activation enables you to avoid that problem. You can also use local service profile activation to provide default services when RADIUS servers are down.

You can apply services to all subscribers in a tunnel group or to all subscribers using a particular LAC. You can configure a maximum of 12 services per tunnel group or LAC hostname.

After configuring one or more dynamic service profiles that define services, you apply them in the tunnel group or in the access profile configuration for a LAC client by specifying the service profile names. You can list more than one profile to be activated, separated by an ampersand (&). You can also specify parameters to be used by the service profile that might override values configured in the profile itself, such as a downstream shaping rate for a CoS service.

The locally configured list of services (via service profiles) serves as local authorization that is applied by authd during client session activation. This list of services is subject to the same validation and processing as services originating from external authority, such as RADIUS. These services are presented during subscriber login.

You can still use RADIUS VSAs or CoA requests in concert with the service profiles. If services are sourced from an external authority as authorization during authentication or during subscriber session provisioning (activation), the services from the external authority take strict priority over those in the local configuration. If a service applied with RADIUS is the same as a service applied with a service profile in the CLI, but with different parameters, the RADIUS service is applied with a new session ID and takes precedence over the earlier service profile.

You can issue commands to deactivate or reactivate any service you have previously activated for a tunnel group or LAC.

Define the dynamic service profiles that you want to later apply to a tunnel group or LAC.

To apply service profiles to all subscribers in a tunnel group:

- Specify one or more service profiles and any parameters to be passed to the services.

```
[edit services l2tp tunnel-group group-name]
user@host# set service-profile profile-name(parameter)&profile-name
```

To apply service profiles to all subscribers for a particular LAC:

- Specify one or more service profiles and any parameters to be passed to the services.

```
[edit access profile profile-name client client-name l2tp]
user@host# set service-profile profile-name (parameter)&profile-name
```



NOTE: When service profiles are configured for a LAC client and for a tunnel group that uses that client, only the LAC client service profile is applied. It overrides the tunnel group configuration. For example, in the following configuration, the tunnel group, tg-LAC-3, uses the LAC client, LAC-3, so the LAC3 configuration overrides the tunnel group configuration. Consequently only the cos-A3 service is activated for subscribers in the tunnel group, rather than Cos2 and fw1. The shaping rate passed for the service is 24 Mbps.

```
[edit]
user@host# set services l2tp tunnel-group tg-LAC-3 service-profile
cos2(31000000)&fw1
user@host# set access profile prof-lac client LAC-3 l2tp service-profile
cos-A3(24000000)
```

You can deactivate any service applied to a subscriber session by issuing the following command:

```
user@host> request network-access aaa subscriber delete session-id subscriber-session-id
service-profile profile-name
```

You can reactivate any service applied to a subscriber session by issuing the following command:

```
user@host> request network-access aaa subscriber add session-id subscriber-session-id
service-profile profile-name
```

To display the services sessions for all current subscriber sessions, use the **show subscribers extensive** or **show network-access aaa subscribers session-id *id-number* detail** command.

To understand how local service application works, the following examples illustrate the various configuration possibilities. First, consider the following dynamic service profile configurations, cos2 and fw1:

```
dynamic-profiles {
  cos2 {
    variables {
      shaping-rate default-value 10m;
      shaping-rate-in default-value 10m;
      data-in-filter uid;
      data-in-policer uid;
    }
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          family inet;
        }
      }
    }
  }
}
```

```

    }
  }
}
class-of-service {
  traffic-control-profiles {
    TrafficShaper {
      scheduler-map a;
      shaping-rate "$shaping-rate";
    }
  }
  interfaces {
    "$junos-interface-ifd-name" {
      unit "$junos-interface-unit" {
        output-traffic-control-profile TrafficShaper;
      }
    }
  }
}
}
|
dynamic-profiles {
  fw1 {
    variables {
      v6input default-value v6ingress;
      v6output default-value v6egress;
      input default-value upstrm-filter;
      output default-value dwnstrm-filter;
    }
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          family inet;
        }
      }
    }
  }
}
}

```

The following statement applies both services to all subscribers in tunnel group tg1; a parameter value of 31 Mbps is passed to the cos2 service:

```
[edit]
user@host# set services l2tp tunnel-group tg1 service-profile cos2(31000000)&fw1
```

In the cos2 service profile, the shaping rate is provided by a user-defined variable with a default value of 10m, or 1Mbps. After the L2TP session is up, cos2 and fw1 are activated with service session IDs of 34 and 35, respectively.

```
user@host1> show subscribers extensive
...
```

```

Service Session ID: 34
Service Session Name: cos2
State: Active
Family: inet
Service Activation time: 2018-02-15 15:44:16 IST

```

```
Service Session ID: 35
Service Session Name: fw1
State: Active
Family: inet
Service Activation time: 2018-02-15 15:44:16 IST
Dynamic configuration:
  input: upstrm-filter
  output: dwnstrm-filter
  v6input: v6ingress
  v6output: v6egress
```

The parameter passed to cos2 is used as the value for \$shaping-rate; consequently the shaping rate for the service is adjusted from the default value of 10 Mbps to 31 Mbps, as shown in the following command output. Although the output indicates the adjusting application is RADIUS CoA, the adjustment is a consequence of the parameter passed to the service profile. That operation uses the same internal framework as a CoA and is reported as such.

```
user@host1> show class-of-service interface si-1/0/0.3221225492
Logical interface: si-1/0/0.3221225492, Index: 3221225492

```

Object	Name	Type	Index
Traffic-control-profile	subscriber-tcp-2	Output	23571
Scheduler-map	a	Output	4294967354
Classifier	dscp-ipv6-compatibility	dscp-ipv6	9
Classifier	ipprec-compatibility	ip	13

Adjusting application: RADIUS CoA

```
Adjustment type: absolute
configured-shaping-rate: 31000000
adjustment-value: 31000000
Adjustment overhead-accounting mode: frame mode
Adjustment overhead bytes: 0
Adjustment target: node
Adjustment priority: 1
```

Now the cos2 service is deactivated from the CLI for subscriber session 27.

```
user@host1> request network-access aaa subscriber delete service-profile cos2
session-id 27
Successful completion
```

The following output shows cos2 is gone, leaving only fw1 as an active service.

```
user@host1> show subscribers extensive
Type: L2TP
User Name: user@example.com
IP Address: 192.0.2.103
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: si-1/0/0.3221225492
Interface type: Dynamic
Underlying Interface: si-1/0/0.3221225492
Dynamic Profile Name: dyn-Ins-profile
State: Active
Radius Accounting ID: 27
```

```

Session ID: 27
PFE Flow ID: 42
Login Time: 2017-08-30 07:29:39 IST
Service Sessions: 1
IP Address Pool: ipv4_pool
Accounting interval: 600
Frame/cell mode: Frame
Overhead accounting bytes: -38
Calculated downstream data rate: 1000000 kbps
Adjusted downstream data rate: 1000000 kbps

```

```

Service Session ID: 35
Service Session Name: fw1
State: Active
Family: inet
Service Activation time: 2018-02-15 15:44:16 IST
Dynamic configuration:
  input: upstrm-filter
  output: dwnstrm-filter
  v6input: v6ingress
  v6output: v6egress

```

The following command reactivates cos2 for subscriber session 27.

```

user@host1> request network-access aaa subscriber add service-profile cos2
session-id 27
Successful completion

```

The reactivated cos2 service has a new service session ID of 36.

```

user@host1> show subscribers extensive
...
Service Session ID: 35
Service Session Name: fw1
State: Active
Family: inet
Service Activation time: 2018-02-15 15:44:16 IST
Dynamic configuration:
  input: upstrm-filter
  output: dwnstrm-filter
  v6input: v6ingress
  v6output: v6egress

Service Session ID: 36
Service Session Name: cos2
State: Active
Family: inet
Service Activation time: 2018-02-15 15:58:23 IST

```

The reactivated cos2 service uses the default shaping rate, 10 Mbps, from the service profile.

```

user@host1> show class-of-service interface si-1/0/0.3221225492
Logical interface: si-1/0/0.3221225492, Index: 3221225492

```

Object	Name	Type	Index
Traffic-control-profile	subscriber-tcp-2	Output	23571
Scheduler-map	a	Output	4294967354
Classifier	dscp-ipv6-compatibility	dscp-ipv6	9

Classifier ipprec-compatibility ip 13

Adjusting application: RADIUS CoA

Adjustment type: absolute
 configured-shaping-rate: 10000000
adjustment-value: 10000000
 Adjustment overhead-accounting mode: frame mode
 Adjustment overhead bytes: 0
 Adjustment target: node
 Adjustment priority: 1

Next, a RADIUS CoA request is received, which includes the Activate-Service VSA (26-65). The VSA specifies and activates the service and specifies a change in the shaping rate of cos2 from the default 10 Mbps to 12 Mbps. The cos2 service session 36 still appears in the output, but is superseded by the new service session initiated by the CoA, 49.

user@host1> show subscribers extensive

...

Service Session ID: 35

Service Session Name: fw1

State: Active
 Family: inet
 Service Activation time: 2018-02-15 15:44:16 IST
 Dynamic configuration:
 input: upstrm-filter
 output: dwnstrm-filter
 v6input: v6ingress
 v6output: v6egress

Service Session ID: 36

Service Session Name: cos2

State: Active
 Family: inet
 Service Activation time: 2018-02-15 15:58:23 IST

Service Session ID: 49

Service Session Name: cos2

State: Active
 Family: inet
 Service Activation time: 2018-02-15 16:25:04 IST
 Dynamic configuration:
 shaping-rate: 12000000
 shaping-rate-in: 10m

user@host1> show class-of-service interface si-1/0/0.3221225492

Logical interface: si-1/0/0.3221225492, Index: 3221225492

Object	Name	Type	Index
Traffic-control-profile	subscriber-tcp-2	Output	23571
Scheduler-map	a	Output	4294967354
Classifier	dscp-ipv6-compatibility	dscp-ipv6	9
Classifier	ipprec-compatibility	ip	13

Adjusting application: RADIUS CoA

Adjustment type: absolute
 configured-shaping-rate: 12000000
adjustment-value: 12000000
 Adjustment overhead-accounting mode: frame mode
 Adjustment overhead bytes: 0

```
Adjustment target: node
Adjustment priority: 1
```

When a service is applied by both the CLI configuration and a RADIUS VSA (26-65), but with different parameters, the RADIUS configuration overrides the CLI configuration. In the following example, the CLI configuration applies the cos2 service profile with a value of 31 Mbps for the shaping rate.

```
[edit]
user@host# set services l2tp tunnel-group tg1 service-profile cos2(31000000)
```

The RADIUS Access-Accept message service activation VSA (26-65) applies cos2 with a value of 21 Mbps for the shaping rate.

```
l2tp@l2tp.com  User-Password := "bras"
                Auth-Type = Local,
                Service-Type = Framed-User,
                Framed-Protocol = PPP,
                ERX-Service-Activate:1 += 'cos2(21000000)',
```

The CLI configuration activates service session 22 with a shaping rate of 31 Mbps. The RADIUS VSA activates service session 23 with a shaping rate of 21 Mbps.

```
user@host1> show subscribers extensive
```

```
...
Service Session ID: 22
Service Session Name: cos2
State: Active
Family: inet
Service Activation time: 2018-02-16 08:22:03 IST
Dynamic configuration:
  shaping-rate: 31000000
  shaping-rate-in: 10m

Service Session ID: 23
Service Session Name: cos2
State: Active
Family: inet
Service Activation time: 2018-02-16 08:22:03 IST
Dynamic configuration:
  shaping-rate: 21000000
  shaping-rate-in: 10m
```

```
user@host1> show class-of-service interface si-1/0/0.3221225492
Logical interface: si-1/0/0.3221225492, Index: 3221225492
```

Object	Name	Type	Index
Traffic-control-profile	subscriber-tcp-2	Output	23571
Scheduler-map	a	Output	4294967354
Classifier	dscp-ipv6-compatibility	dscp-ipv6	9
Classifier	ipprec-compatibility	ip	13

Adjusting application: RADIUS CoA

```
Adjustment type: absolute
configured-shaping-rate: 21000000
adjustment-value: 21000000
Adjustment overhead-accounting mode: frame mode
Adjustment overhead bytes: 0
Adjustment target: node
Adjustment priority: 1
```

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, you can apply services to L2TP sessions by means of dynamic service profiles without involving RADIUS.

Related Documentation

- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)
- *Activating and Deactivating Subscriber Services Locally with the CLI*

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions

You can create a pool of inline service interfaces, also known as a *service device pool*, to enable load-balancing of L2TP traffic across the interfaces. The pool is supported for dynamic LNS configurations, where it provides a set of logical interfaces that can be dynamically created and allocated to L2TP sessions on the LNS. The pool is assigned to an LNS tunnel group. L2TP maintains the state of each inline service interface and uses a round-robin method to evenly distribute the load among available interfaces when new session requests are accepted.



NOTE: Load balancing is available only for dynamically created subscriber interfaces.

LNS sessions anchored on an MPC are not affected by a MIC failure as long as some other path to the peer LACs exists. If the MPC hosting the peer interface fails and there is no path to peer LACs, the failure initiates termination and clean-up of all the sessions on the MPC.

If the MPC anchoring the LNS sessions itself fails, the Routing Engine does not relocate sessions to another slot and all sessions are terminated immediately. New sessions can come up on another available interface when the client retries.

To configure the service device pool:

1. Create the pool.

```
[edit services service-device-pools]
user@host# edit pool pool-name
```

2. Specify the inline service interfaces that make up the pool.

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


- Related Documentation**
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
 - [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)

Configuring a Dynamic Profile for Dynamic LNS Sessions

You can configure L2TP to dynamically assign inline service interfaces for L2TP tunnels. You must define one or more dynamic profiles and assign a profile to each tunnel group. The LNS supports IPv4-only, IPv6-only, and dual-stack IPv4/IPv6 sessions.

To configure the L2TP dynamic profile:

1. Create the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Configure the interface to be dynamically assigned to the routing instance used by the tunneled PPP clients.

```
[edit dynamic-profiles profile-name routing-instances "$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

3. Configure the routing options for access routes in the routing instance.

```
[edit dynamic-profiles profile-name routing-instances "$junos-routing-instance"
routing-options access]
user@host# set route next-hop $junos-framed-route-nexthop
user@host# set route metric $junos-framed-route-cost
user@host# set route preference $junos-framed-route-distance
```

4. Configure the routing options for access-internal routes in the routing instance.

```
[edit dynamic-profiles profile-name routing-instances "$junos-routing-instance"
routing-options access-internal]
user@host# set route $junos-subscriber-ip-address
```

5. Define the interfaces used by the dynamic profile. The variable is dynamically replaced by one of the configured inline service interfaces.

```
[edit dynamic-profiles profile-name]
user@host# set interfaces $junos-interface-ifd-name
```

6. Configure the inline services logical interfaces to be dynamically instantiated.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name"]
user@host# set unit $junos-interface-unit
```

7. Specify an identifier for the logical interfaces.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
```

```
user@host# set dial-options l2tp-interface-id name
```

8. Configure each logical interface to be used for only one session at a time.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set dial-options dedicated
```

9. Configure the address family for the logical interfaces and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.



NOTE: Dynamic LNS sessions require you to include the `dial-options` statement in the dynamic profile, which in turn requires you to include the `family inet` statement. This has the following consequences:

- You must always configure `family inet` regardless of whether you configure IPv4-only, IPv6-only, or dual-stack interfaces in the profile.
- When you configure IPv4-only interfaces, you configure only `family inet` and you must configure the interface address under `family inet`.
- When you configure IPv6-only interfaces, you must also configure `family inet6` and you must configure the interface address under `family inet6`. You do not configure the address under `family inet`.
- When you configure dual-stack, IPv4/IPv6 interfaces, you configure both `family inet` and `family inet6` and an interface address under each family.

For IPv4-only interfaces:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set family inet unnumbered-address $junos-loopback-interface
```

For IPv6-only interfaces:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set family inet  
user@host# set family inet6 unnumbered-address $junos-loopback-interface
```

For dual-stack IPv4/IPv6 interfaces:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit"]  
user@host# set family inet unnumbered-address $junos-loopback-interface  
user@host# set family inet6 unnumbered-address $junos-loopback-interface
```



.....

NOTE: If Router Advertisement Protocol is configured, then you configure a numbered address rather than an unnumbered address for the IPv6 local address:

```
user@host# set family inet6 address $junos-ipv6-address
```

See [Broadband Subscriber Sessions Feature Guide](#) for information about using variables for IPv6-only and dual-stack addressing in dynamic profiles.

.....

**Related
Documentation**

- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)

CHAPTER 24

Configuring IP Packet Fragment Reassembly

- [IP Packet Fragment Reassembly for L2TP Overview on page 289](#)
- [Configuring IP Inline Reassembly for L2TP on page 290](#)

IP Packet Fragment Reassembly for L2TP Overview

You can configure the service interfaces on the MX Series routers with modular port concentrators (MPCs) to support reassemble fragmented IP packets for an L2TP connection. When packets are transmitted over an L2TP connection, the packets may be fragmented during transmission and need to be reassembled before they are processed further. Efficient reassembly is important for network throughput, scalability, and graceful response to congestion.

Fragmentation of IP packets for transmission and the need to reassemble the IP packets at a destination is a feature of how Layer 2 (the frame layer) and Layer 3 (the packet layer) operate. The maximum size of a frame, set by the Maximum Transmission Unit (MTU) value, and the maximum size of a packet are determined independently. Typically the packet size can far exceed the MTU size defined for the outgoing connection. If the packet size (data plus IP and other headers) exceeds the configured frame size (usually set by the transport medium limits), the packet must be fragmented and split across multiple frames for transmission. Frames are always processed immediately, when they arrive (if error-free), but packet fragments cannot be processed until the whole packet has been reassembled. Each packet fragment inside a frame series, except the last packet fragment, has the more fragments (MF) IP header bit set, indicating that this packet is part of a whole. The last packet fragment inside a frame does not have this MF bit set and therefore ends the fragment sequence. After all of the fragments of a packet have arrived, the entire packet can be reassembled.

In an L2TP connection, packets are transmitted between the L2TP Access Concentrator (LAC) and the L2TP Network Server (LNS). For an IP packet being transmitted over an L2TP connection, the packet is fragmented at the LAC, at an LNS, or at any intermediate router. IP reassembly parameters configured on the service interfaces of the LAC and the LNS determine how the fragments are reassembled at the service interfaces to ensure efficient reassembly over an L2TP connection.

- Related Documentation**
- [Configuring IP Inline Reassembly for L2TP on page 290](#)
 - [Protocols and Applications Supported on the MPC1E for MX Series Routers](#)
 - [ip-reassembly on page 514](#)

Configuring IP Inline Reassembly for L2TP

This procedure shows how to configure a service interface on a LAC or LNS to reassemble fragmented IP packets. This example creates a service set that configures the IP reassembly parameters for L2TP fragments. The service set is then associated with the L2TP service.

Before you configure inline IP reassembly, be sure you have:

- Configured L2TP.
- Configured a valid service interface on the LAC or LNS.

To configure inline IP reassembly:

1. Configure the chassis-level bandwidth used by the inline services interface on the FPC and PIC slot for inline IP fragment reassembly.

```
[edit chassis]
user@host# set fpc 2 pic 1 inline-services bandwidth 10g
```

2. Configure the interface-level logical unit used by the inline services (si-) interface on the FPC and PIC slot for inline IP fragment reassembly.

```
[edit interfaces]
user@host# set si-2/1/0 unit 0 family inet
user@host# set si-2/1/0 unit 0 service-domain inside
```



NOTE: This configuration is not unique to L2TP. However, you must configure the family (inet) and service domain (inside) as shown.

3. Configure the service set (**set1**) for IP reassembly in the input match direction. (The **local** option loops the reassembled packets back to the local interface.)

```
[edit services]
user@host# set service-set set1
```

```
[edit services service-set set1]
user@host# set ip-reassembly-rules ipr_rule1
user@host# set next-hop-service inside-service-interface si-2/1/0.0
user@host# set next-hop-service outside-service-interface-type local
```

**NOTE:**

- You must configure both inside (si- interface) and outside type (local) service interfaces statements. The reassembly rule is not formulated outside of the service set; this statement simply initiates the reassembly process.
- You can configure only one service interface for each service-set.

4. Configure the IP reassembly rule parameter.

```
[edit services ip-reassembly]
user@host# set rule ipr_rule1 match-direction input;
```

5. Configure the service set (**set1**) for IP reassembly to bind to the L2TP service.

**NOTE:**

- The service set must be defined at the [edit services] hierarchy level.
- You cannot delete a service set instance if it is associated with an L2TP service.

```
[edit services l2tp]
user@host# set ip-reassembly service-set set1
```

**Related
Documentation**

- [IP Packet Fragment Reassembly for L2TP Overview on page 289](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- *Protocols and Applications Supported on the MPC1E for MX Series Routers*

Configuring High Availability in the L2TP Access Network

- [L2TP and Graceful Routing Engine Switchover on page 293](#)
- [L2TP Failover and Peer Resynchronization on page 294](#)
- [Configuring the L2TP Peer Resynchronization Method on page 295](#)
- [High Availability Using Unified ISSU in the L2TP Access Network on page 297](#)
- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 298](#)

L2TP and Graceful Routing Engine Switchover

Graceful Routing Engine switchover (GRES) is supported on MX Series routers acting as either the L2TP LAC or LNS. In the event that L2TP (the l2tp-universal-edge process) restarts or that the router fails over from the active routing engine (RE) to the standby RE, L2TP graceful Routing Engine switchover ensures that the following occurs:

- The LAC and the LNS recover destinations, tunnels, and sessions that were already established at the time of the failure or restart.
- The LAC and the LNS respond to tunnel keepalive requests received during the switchover for established tunnels, but do not generate any keepalives until the switchover is complete.
- The LAC and the LNS delete all the tunnels and sessions that are not in the Established state.
- The LAC and the LNS reject requests to create new tunnels and sessions.
- The LAC and the LNS send another disconnect notification to the peer for sessions and tunnels that are already in the Disconnecting state at the time of the failure or restart. For sessions and tunnels that were coming up at that time, the LAC and LNS send a disconnect notification to the peer.
- The LAC and the LNS restart timers for the full timeout period for recovered L2TP destinations, tunnels, and sessions.

If a graceful Routing Engine switchover (GRES) is triggered by an operational mode command, the state of aggregated services interfaces (ASIs) are not preserved. For example:

```
request interface <switchover | revert> asi-interface
```

However, if GRES is triggered by a CLI commit or FPC restart or crash, the backup Routing Engine updates the ASI state. For example:

```
set interface si-x/y/z disable
commit
```

Or:

```
request chassis fpc restart
```

- Related Documentation**
- [L2TP Failover and Peer Resynchronization on page 294](#)
 - [L2TP for Subscriber Access Overview on page 147](#)

L2TP Failover and Peer Resynchronization

L2TP failover enables a failed L2TP endpoint to resynchronize with its nonfailed peer during recovery and restart of the L2TP protocol on the failed endpoint. L2TP failover is enabled by default.

The failover and L2TP peer resynchronization process does all of the following:

- Prevents the nonfailed endpoint from prematurely terminating a tunnel while the failed endpoint is recovering.
- Reestablishes the sequence numbers required for the operation of the L2TP control protocol.
- Resolves inconsistencies in the tunnel and session databases of the failed endpoint and the nonfailed endpoint.

The router supports both the L2TP failover protocol method (described in *RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*) and the L2TP silent failover method. The differences between these two methods are as follows:

- The L2TP failover protocol method requires a nonfailed endpoint to wait an additional recovery time period while the failed endpoint is recovering to prevent the nonfailed endpoint from prematurely disconnecting the tunnel. The additional recovery period delays the detection of tunnel keepalive failures.

If a peer on an MX series router negotiates failover protocol with an MX Series peer that is not configured for failover protocol, both use the silent failover method. If the negotiation is with a third-party device that does not support failover protocol, the MX Series peer falls back to silent failover; whether the third-party peer recovers in this case depends on how resynchronization is implemented on that device.

- Silent failover operates entirely within the failed endpoint and does not require nonfailed endpoint support—this improves interoperability between peers. Silent failover does not require additional recovery time by the nonfailed endpoint, which also eliminates the potential for degraded responsiveness to the loss of tunnel connectivity. Starting in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, silent failover is the default resynchronization method in Junos OS.

In lower-numbered releases, the default resynchronization method is *failover-protocol-fall-back-to-silent-failover*. The recovery method used depends on the results of the failover capability negotiation that takes place between L2TP peers when they establish a tunnel, which works as follows:

- L2TP on the LAC by default attempts to negotiate the L2TP failover protocol first. When L2TP determines that the remote peer supports the L2TP failover protocol, then the L2TP failover protocol method is used.
- When L2TP determines that the remote peer does not support the L2TP failover protocol, then the L2TP silent failover method is used. Falling back on this secondary method prevents the failover from forcing a disconnection of the tunnel to the peer and all its sessions.

In Junos OS releases where *failover-protocol-fall-back-to-silent-failover* is the default method, you can change the default behavior by including the [disable-failover-protocol](#) statement at the **[edit services l2tp]** hierarchy level. This statement forces the configured LAC or LNS endpoint to operate only in silent failover mode. This configuration can be used to prevent the device from negotiating failover protocol with the peer even if the peer tries to negotiate it. When you issue this statement and the peer supports only failover protocol, the nonfailed endpoint (LAC or LNS) uses silent failover for recovery. Starting in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the [disable-failover-protocol](#) statement is deprecated, because the change in default resynchronization method makes it unnecessary.

Release History Table

Release	Description
15.1R6	Starting in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, silent failover is the default resynchronization method in Junos OS.
15.1R6	Starting in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the disable-failover-protocol statement is deprecated, because the change in default resynchronization method makes it unnecessary.

Related Documentation

- [Configuring the L2TP Peer Resynchronization Method on page 295](#)
- [L2TP and Graceful Routing Engine Switchover on page 293](#)
- [L2TP for Subscriber Access Overview on page 147](#)

Configuring the L2TP Peer Resynchronization Method

The L2TP implementation on MX Series routers supports resynchronization between a failed L2TP endpoint and its peer nonfailed endpoint. Peer resynchronization enables L2TP to recover from a daemon or router restart or a Routing Engine switchover.

L2TP peer resynchronization:

- Prevents the nonfailed endpoint from prematurely terminating a tunnel while the failed endpoint is recovering.

- Reestablishes the sequence numbers required for the operation of the L2TP control protocol.
- Resolves inconsistencies in the tunnel and session databases of the failed endpoint and the nonfailed endpoint.

You can configure the peer resynchronization method you want the router to use. Both the L2TP failover protocol method and the L2TP silent failover method are supported.

In Junos OS Releases through 15.1R5, 16.1R4, 16.2R1, and 17.1R1, the default behavior is for L2TP on the LAC to attempt to negotiate the L2TP failover protocol with the LNS. When the LNS supports this method and negotiation is successful, the L2TP failover protocol is used when either peer fails. When negotiation for L2TP failover protocol fails, then the peers use silent failover when either peer fails. This behavior is called failover-protocol-fall-back-to-silent-failover. Falling back to the silent failover method when failover protocol negotiation is unsuccessful prevents a subsequent peer failure from forcing a disconnection of the tunnel to the peer and all the associated sessions.



NOTE: The behavior just described applies when both peers are MX Series routers. If one endpoint is a third-party device, then the behavior for that device depends on its L2TP implementation.

You can disable the default behavior and force the LAC or the LNS to operate only in silent failover mode. This configuration can be useful when the peer routers either are configured for silent failover or incorrectly negotiate to use the failover protocol even though they do not support it. Another reason to use this statement is that the failover protocol method keeps the tunnel open with the failed peer, in case the failed peer is able to recover from the failure and resynchronize with the nonfailed peer. This behavior keeps the tunnel up and the subscribers logged in while traffic is not flowing, preventing service level agreements from being met. When you issue this statement and the peer supports only failover protocol, the nonfailed endpoint (LAC or LNS) uses silent failover for recovery.

To disable negotiation of the L2TP failover protocol:

- Configure disabling.

```
[edit services l2tp]  
user@host# set disable-failover-protocol
```

Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the default failover resynchronization method is changed to silent failover. Consequently, the **disable-failover-protocol** statement no longer needs to be used and is deprecated. If you upgrade from a lower-numbered release where the default method is failover-protocol-fall-back-to-silent-failover, and your configuration includes the **disable-failover-protocol** statement, the configuration is still supported, but the CLI notifies you that the statement is deprecated.

In these releases, you can still configure which method you want an endpoint to use, failover protocol or silent failover.

To configure the LAC or LNS to negotiate the L2TP failover protocol:

- Specify the failover protocol.

```
[edit services l2tp]
user@host# set failover-resync failover-protocol
```

If the negotiation fails, the endpoint falls back to the silent failover method.

To restore the default resynchronization method for the LAC or LNS:

- Specify the silent failover method.

```
[edit services l2tp]
user@host# set failover-resync silent-failover
```

Release History Table

Release	Description
15.1R6	Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the default failover resynchronization method is changed to silent failover. Consequently, the disable-failover-protocol statement no longer needs to be used and is deprecated.
15.1R5	In Junos OS Releases through 15.1R5, 16.1R4, 16.2R1, and 17.1R1, the default behavior is for L2TP on the LAC to attempt to negotiate the L2TP failover protocol with the LNS.

**Related
Documentation**

- [L2TP Failover and Peer Resynchronization on page 294](#)

High Availability Using Unified ISSU in the L2TP Access Network

Starting in Junos OS Release 14.1, the unified in-service software upgrade (unified ISSU) feature supports the L2TP access model used by subscriber management. This support ensures that the router preserves active L2TP subscriber sessions and session services after a unified ISSU has completed.

See *Getting Started with Unified In-Service Software Upgrade* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

The LAC and the LNS support unified ISSU. When an upgrade is initiated, the LAC completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade. Subscriber logouts are recorded during the upgrade and are completed after the upgrade has completed.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, the unified in-service software upgrade (unified ISSU) feature supports the L2TP access model used by subscriber management. This support ensures that the router preserves active L2TP subscriber sessions and session services after a unified ISSU has completed.

Related Documentation

- [Verifying and Monitoring Subscriber Management Unified ISSU State on page 111](#)
- [Unified ISSU System Requirements](#)

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane quiescing has not started). The second example shows that unified ISSU is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
```

```
General:
```

```
Graceful Restart    Enabled
Mastership          Master
Database            Available
Chassisd ISSU State IDLE
ISSU State          IDLE
ISSU Wait           0
```

```
user@host> show system subscriber-management summary
```

```
General:
```

```
Graceful Restart    Enabled
Mastership          Master
Database            Available
Chassisd ISSU State DAEMON_ISSU_PREPARE
ISSU State          PREPARE
ISSU Wait           198
```

Related Documentation

- [High Availability Using Unified ISSU in the PPP Access Network on page 139](#)
- [High Availability Using Unified ISSU in the DHCP Access Network on page 106](#)
- [High Availability Using Unified ISSU in the L2TP Access Network on page 297](#)
- [Getting Started with Unified In-Service Software Upgrade](#)

CHAPTER 26

Monitoring and Managing L2TP for Subscriber Access

- [Verifying and Managing L2TP for Subscriber Access on page 299](#)
- [Testing L2TP Tunnel Configurations from the LAC on page 300](#)
- [Enabling Tunnel and Global Counters for SNMP Statistics Collection on page 303](#)

Verifying and Managing L2TP for Subscriber Access

Purpose View or clear information about L2TP tunnels and sessions.



BEST PRACTICE: The all option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the all option with the clear services l2tp destination, clear services l2tp session, or clear services l2tp tunnel statements in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

- Action**
- To display a summary of L2TP tunnels, sessions, errors, and control and data packets:
user@host> **show services l2tp summary**
 - To display the L2TP destinations:
user@host> **show services l2tp destination**
 - To clear all L2TP destinations:
user@host> **clear services l2tp destination all**
 - To clear statistics for all L2TP tunnels belonging to a destination, tunnels belonging to a specified local-gateway address, and tunnels belonging to a specified peer-gateway address:
user@host> **clear services l2tp destination statistics all**
user@host> **clear services l2tp destination local-gateway 203.0.113.2**
 - To display the L2TP sessions:
user@host> **show services l2tp session**

- To clear all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

```
user@host>clear services l2tp session all
user@host>clear services l2tp session local-session-id 40553
user@host>clear services l2tp session local-gateway 203.0.113.2
user@host>clear services l2tp session local-gateway-name lns-mx960
```

- To clear statistics for all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

```
user@host>clear services l2tp session statistics all
user@host>clear services l2tp session statistics local-session-id 17967
user@host>clear services l2tp session statistics local-gateway 203.0.113.2
user@host>clear services l2tp session statistics local-gateway-name lns-mx960
```

- To display the L2TP tunnels:

```
user@host> show services l2tp tunnel
```

- To clear all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

```
user@host> clear services l2tp tunnel all
user@host>clear services l2tp tunnel local-tunnel-id 40553
user@host>clear services l2tp tunnel local-gateway 203.0.113.2
user@host>clear services l2tp tunnel local-gateway-name lns-mx960
```

- To clear statistics for all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

```
user@host> clear services l2tp tunnel statistics all
user@host>clear services l2tp tunnel statistics local-tunnel-id 40553
user@host>clear services l2tp tunnel statistics local-gateway 203.0.113.2
user@host>clear services l2tp tunnel statistics local-gateway-name lns-mx960
```

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [CLI Explorer](#)

Testing L2TP Tunnel Configurations from the LAC

You can test L2TP tunnel configurations on the LAC and successful subscriber authentication and tunneling without bringing up a PPP user and an associated tunnel.

Issue the **test services l2tp tunnel** command from CLI operational mode to map a subscriber to an L2TP tunnel, verify the L2TP tunnel configuration (both locally on the LAC and on a back-end server such as a RADIUS server), and verify that L2TP tunnels from the LAC can be established with the remote LNS.

The Junos OS LAC implementation enables you to configure multiple tunnels from which one tunnel is chosen for tunneling a PPP subscriber. You can use the **test services l2tp**

tunnel command to test all possible tunnel configurations to verify that each can be established. Alternatively, you can test only a specific tunnel for the subscriber.

You must specify a configured subscriber username when you issue the command. The test generates a dummy password—*testpass*—for the subscriber, or you can optionally specify the password. The test verifies whether the subscriber identified by that username can be tunneled according to the tunnel configuration. If the subscriber can be tunneled, then the test verifies whether the L2TP tunnel can be established with the LNS according to the L2TP configuration.

You can optionally specify a tunnel ID, in which case only that tunnel is tested; the tunnel must be already configured for that username. If you omit this option, the test is applied to the full set of tunnel configurations that are returned for the username. The tunnel ID you specify is the same as that used by Tunnel-Assignment-Id (RADIUS attribute 82) and specified by the **identification** statement in the tunnel profile.

To test subscriber authentication and tunnel configuration:

- Specify only the username.

Example 1:

```
user@host> test services l2tp tunnel user test-user1@example.com
Subscriber: test-user1@example.com, authentication failed
```

The user failed authentication with the generated password and consequently was not tunneled.

Example 2:

```
user@host> test services l2tp tunnel user user23@example.com
Subscriber: user23@example.com, authentication success, l2tp tunneled
```

Tunnel-name	Tunnel-peer	Logical-System	Routing-Instance	Status
test1tunnel	192.168.2.3	default	default	Up
test2tunnel	192.168.30.3	default	default	Peer unresponsive
test3tunnel	192.168.50.1	default	test	Up

This user was authenticated with the generated password and successfully tunneled. A set of tunnels was found to be associated with that username and the entire set was tested.

- Specify the username and the user's configured password.

```
user@host> test services l2tp tunnel user test-user1@example.com password grZ98#jW
Subscriber: test-user1@example.com, authentication success, locally terminated
```

The subscriber was authenticated. However, the user was terminated locally rather than tunneled; this means that no tunnel was found to be associated with the user.

- Specify the username and a particular tunnel for the subscriber.

```
user@host> test services l2tp tunnel user rx37w@example.com tunnel-name ce-lac
Subscriber: rx37w@example.com, authentication success, l2tp tunneled
```

Tunnel-name	Tunnel-peer	Logical-System	Routing-Instance	Status
ce-lac	192.168.5.10	default	default	Up

The subscriber was authenticated and tunneled. The specified tunnel was found for the subscriber and the tunnel was established, confirming the tunnel configuration.

- Specify the username, the user's configured password, and a tunnel.

```
user@host> test services l2tp tunnel user fanta4-mfg-fan@example.com password dieda499
tunnel-name tunnel5
Subscriber: fanta4-mfg-fan@example.com, authentication success, l2tp tunneled
```

The subscriber was authenticated and tunneled. The absence of tunnel information in the output indicates that the specified tunnel configuration does not exist.

**Related
Documentation**

- [L2TP for Subscriber Access Overview on page 147](#)

Enabling Tunnel and Global Counters for SNMP Statistics Collection

By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in [Table 25 on page 303](#) have a default value of zero.

Table 25: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 60,000 sessions, none of these statistics can be more than 30 minutes old.



BEST PRACTICE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

To enable L2TP statistics collection for SNMP:

- Enable statistics collection.

```
[edit services l2tp]
user@host1# set enable-snmp-tunnel-statistics
```

Related Documentation

- [Configuring an L2TP LAC on page 173](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

PART 5

Configuring MPLS Pseudowire for Subscribers

- [Configuring MPLS Pseudowire Subscriber Logical Interfaces on page 307](#)
- [Configuring Hierarchical CoS Pseudowire Subscriber Interfaces on page 323](#)
- [Configuring CoS Two-Level Hierarchical Scheduling on page 327](#)
- [Configuring CoS Three-Level Hierarchical Scheduling on page 331](#)

CHAPTER 27

Configuring MPLS Pseudowire Subscriber Logical Interfaces

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview on page 310](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 315](#)
- [Configuring a Pseudowire Subscriber Logical Interface Device on page 315](#)
- [Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface on page 318](#)
- [Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces on page 318](#)
- [Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces on page 319](#)
- [Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface on page 321](#)

Pseudowire Subscriber Logical Interfaces Overview

Subscriber management supports the creation of subscriber interfaces over point-to-point MPLS pseudowires. The pseudowire subscriber interface capability enables service providers to extend an MPLS domain from the access-aggregation network to the service edge, where subscriber management is performed. Service providers can take advantage of MPLS capabilities such as failover, rerouting, and uniform MPLS label provisioning, while using a single pseudowire to service a large number of DHCP and PPPoE subscribers in the service network.




NOTE: Pseudowire subscriber logical interfaces are supported on Modular Port Concentrators (MPCs) with Ethernet Modular Interface Cards (MICs) only.

The pseudowire is a tunnel that is either an MPLS-based Layer 2 VPN or Layer 2 circuit. The pseudowire tunnel transports Ethernet encapsulated traffic from an access node

(for example, a DSLAM or other aggregation device) to the MX Series router that hosts the subscriber management services. The termination of the pseudowire tunnel on the MX Series router is similar to a physical Ethernet termination, and is the point at which subscriber management functions are performed. A service provider can configure multiple pseudowires on a per-DSLAM basis and then provision support for a large number of subscribers on a specific pseudowire. [Figure 14 on page 308](#) shows an MPLS network that provides subscriber management support.


At the access node end of the pseudowire, the subscriber traffic can be groomed into the pseudowire in a variety of ways, limited only by the number and types of interfaces that can be stacked on the pseudowire. You specify an anchor point, which identifies the logical tunnel interface that terminates the pseudowire tunnel at the access node.

Figure 14: MPLS Access Network with Subscriber Management Support



[Figure 15 on page 309](#) shows the protocol stack for a pseudowire subscriber logical interface. The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD), and supports a circuit-oriented Layer 2 protocol (either Layer 2 VPN or Layer 2 circuit). The Layer 2 protocol provides the transport and service logical interfaces, and supports the protocol family (IPv4, IPv6, or PPPoE).

Figure 15: Pseudowire Subscriber Interface Protocol Stack



The pseudowire configuration is transparent to the subscriber management applications and has no impact on the packet payloads that are used for subscriber management. Subscriber applications such as DHCP and PPPoE can be stacked over Layer 2 similar to the way in which they are stacked over a physical interface.



NOTE: 1. Starting with Junos OS release 16.1, family inet and family inet6 are supported on the services side of an MPLS pseudowire subscriber as well as non-subscriber logical interface.

2. Starting with Junos OS Release 16.1R1, Inline IPFIX is supported on the services side of an MPLS pseudowire subscriber logical interface.

3. Starting with Junos OS Release 15.1R3 and 16.1R1 and later releases, CCC encapsulation is supported on the transport side of an MPLS pseudowire subscriber logical interface.

4. Starting with Junos OS Release 15.1R3 and 16.1R1 and later releases, distributed denial-of-service (DDoS) protection is supported on the services side of an MPLS pseudowire subscriber logical interface.

5. Starting with Junos OS Release 15.1R3 and 16.1R1 and later releases, Policer and Filter are supported on the services side of an MPLS pseudowire subscriber logical interface.

6. Starting with Junos OS Release 15.1R3 and 16.1R1 and later releases, accurate transmit statistics on logical interface are supported on the services side of an MPLS pseudowire subscriber logical interface.

7. Starting with Junos OS Release 17.3R1 and later releases, stateful anchor point redundancy support is provided for pseudowire subscriber logical interface by the underlying redundant logical tunnel interface (rlt) in active-backup mode. This redundancy protects the access and the core facing link against anchor PFE (Packet Forwarding Engine) failure.

**Related
Documentation**

- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)

Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview

In MPLS pseudowire deployments that use pseudowire subscriber logical interfaces, failure of the Packet Forwarding Engine hosting the logical tunnel that anchors those logical interfaces leads to traffic loss and subsequent subscriber session loss.

The Packet Forwarding Engine does not rely on the control plane for failure detection; instead it uses a liveness detection mechanism, with an underlying heartbeat-based algorithm, to detect the failure of other Packet Forwarding Engines in the system. The failure of a Packet Forwarding Engine also indicates the failure of the hosted logical tunnel, which ultimately lead to session loss. To avoid this session loss, a redundant anchor point is required to which the session can be moved without losing any traffic.

Starting from Junos OS Release 17.3 onward, pseudowire subscriber logical interfaces can be instantiated over an underlying redundant logical tunnel (rlt) interface in active-backup mode. This is in addition to installing pseudowires over a single logical

tunnel interfaces. The most noticeable advantage of implementing the pseudowire subscriber logical interface over redundant logical tunnel interfaces is to provide redundancy of the underlying forwarding path.

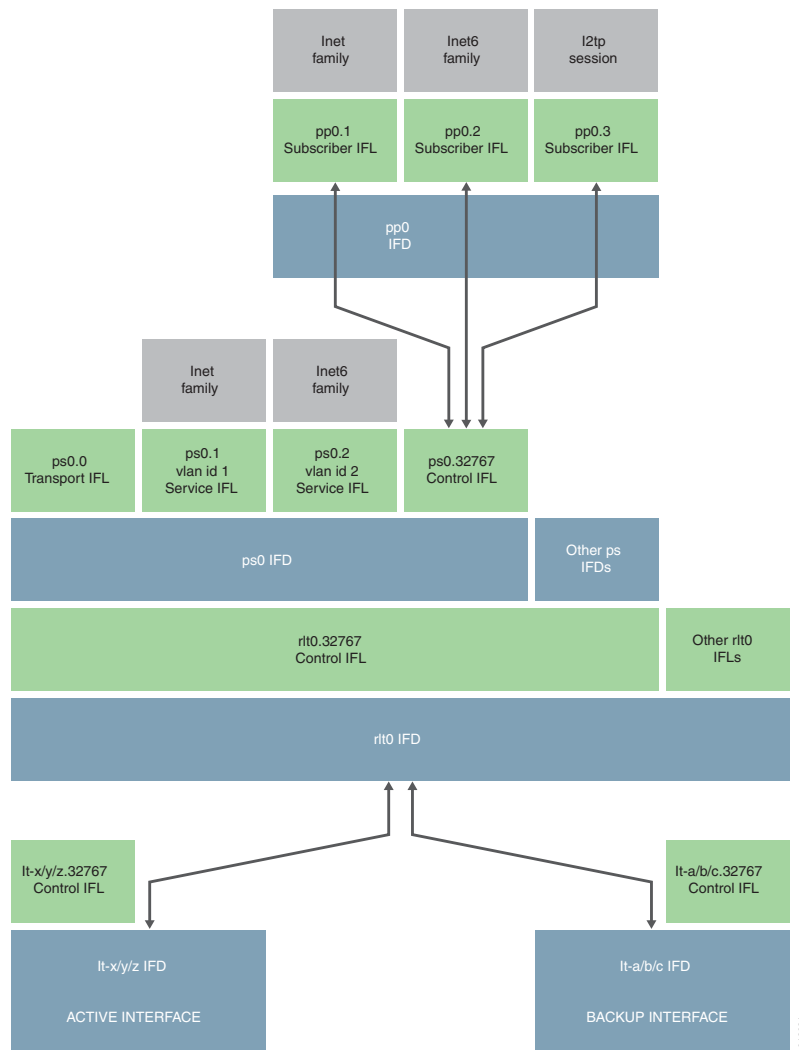
Junos OS Release 17.3 also supports an enhanced aggregated infrastructure for a Packet Forwarding Engine to provide anchor point redundancy. Enhanced aggregated infrastructure requires a minimum of one control logical interface that needs to be created on a redundant logical tunnel interface. Both transport and services logical interfaces created for the pseudowire subscriber logical interface are stacked on the underlying control logical interface for the redundant logical tunnel. This stacking model is used for both redundant and nonredundant pseudowire subscriber logical interfaces.

The following events have to trigger the removal of the physical interface from a redundant group:

- Hardware failure on Modular PIC Concentrator (MPC) or Modular Interfaces Card (MIC).
- MPC failure due to microkernel crash.
- MPC or MIC taken offline administratively.
- Power failure on an MPC or a MIC.

[Figure 16 on page 312](#) provides the details of pseudowire subscriber logical interface stacking over a redundant logical tunnel interface.

Figure 16: Pseudowire Subscriber Logical Interface Stacking over Redundant Logical Tunnel Interface



By default, *Link Protection* for redundant tunnel interfaces is revertive. In case of the active link failure, traffic is routed through the backup link. When the active link is reestablished, traffic is automatically routed back to the active link. This causes traffic loss and subscriber session loss.

To overcome the traffic and session loss, you can configure nonrevertive link protection for redundant tunnel interfaces by using the configuration statement **set interfaces rltX logical-tunnel-options link-protection non-revertive**. With this configuration, when the active link is reestablished, traffic is not routed back to the active link and continue to be forwarded on the backup link. Therefore, there is no traffic loss or subscriber session loss. You can also manually switch traffic from the backup link to the active link by using the **request interface (revert | switchover) interface-name** command.



CAUTION: The manual switching of the traffic incurs traffic loss.



NOTE:

- A control logical interface is created implicitly on a redundant tunnel interface with the pseudowire subscriber logical interface configuration and thus no additional configuration is needed.
- A redundant logical tunnel interface allows 32 member logical tunnel physical interfaces. However, a pseudowire subscriber logical interface hosted on the redundant logical tunnel interface limits the number of logical tunnel physical interfaces to two.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface Device on page 315](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Configuring a Pseudowire Subscriber Logical Interface

A pseudowire subscriber logical interface terminates an MPLS pseudowire tunnel from an access node to the MX Series router that hosts subscriber management, and enables you to perform subscriber management services at the interface.

To create a pseudowire subscriber logical interface:

1. Specify the number of pseudowire logical interfaces that the router can support.
See [“Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router” on page 315](#).
2. Configure the pseudowire subscriber logical interface device.
See [“Configuring a Pseudowire Subscriber Logical Interface Device” on page 315](#).
3. Configure the transport logical interface.
See [“Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface” on page 318](#).
4. Configure the signaling for the pseudowire subscriber interface. You can use either Layer 2 circuit signaling or Layer 2 VPN signaling. The two signaling types are mutually exclusive for a given pseudowire.
 - To configure Layer 2 circuit signaling, see [“Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces” on page 318](#).

- To configure Layer 2 VPN signaling, see [“Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces”](#) on page 319.

5. Configure the service logical interface.

See [“Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface”](#) on page 321.

6. Configure the underlying interface device.

See *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*.

7. Configure CoS parameters and BA classification.

See [“CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces”](#) on page 324.

8. (Optional) Associate a dynamic profile with the pseudowire subscriber logical interface.

You can associate DHCP, PPPoE, IP demux, and VLAN dynamic profiles with pseudowire subscriber logical interfaces. The support is similar to the typical Ethernet interface support.



NOTE: When using a PPPoE dynamic profile to create a pseudowire subscriber logical interface over a demux interface device, the dynamic profile must explicitly specify the correct pseudowire interface device over which the interface is created. The dynamic profile does not automatically create the interface over the demux0 interface device, as is the case with a VLAN demux interface.

9. (Optional) Configure interface set support for pseudowire subscriber logical interfaces.

See *Configuring Interface Sets and Understanding Interface Sets*.

10. (Optional) Stack PPPoE logical interfaces over a pseudowire logical device.

**Related
Documentation**

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)

Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router

You must set the maximum number of pseudowire logical interface devices (pseudowire tunnels) that the router can use for subscriber logical interfaces. You can specify a maximum of 2048 pseudowire logical interface devices for an MX Series router or PTX Series router.

A PFE can host a maximum of 2048 pseudowire logical interface devices, which is the chassis maximum. This PFE hosting support provides the configuration flexibility needed for special cases that might occur for business edge scenarios. However, you can exceed the available PFE resources as you configure additional services on the pseudowire logical interface devices ports. To support a scaled configuration, ensure that you populate the appropriate number of PFEs for the chassis, and that you distribute the pseudowire logical interface devices across the PFEs in such a way that ensures that no PFE is overwhelmed by the anticipated peak load. As part of the network planning for your particular deployment, you must consider the exact mix of the distribution of the pseudowire logical interface devices and the services associated with the devices.



BEST PRACTICE: A configured pseudowire logical interface device consumes resources from shared pools even when the device has no active subscriber logical interfaces. To conserve resources, do not deploy an excessive number of pseudowire devices that you do not intend to use.

To configure the number of pseudowire logical interface devices that you want the router to support:

1. Specify that you want to configure the pseudowire service.

```
[edit chassis]
user@host# edit pseudowire-service
```

2. Set the maximum number of pseudowire logical interface devices.

```
[edit chassis pseudowire-service]
user@host# set device-count 500
```

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Configuring a Pseudowire Subscriber Logical Interface Device

To configure a pseudowire logical interface device that the router uses for subscriber logical interfaces, you specify the logical tunnel that processes the pseudowire termination. You can also configure additional optional parameters for the interface device, such as VLAN tagging method, MTU, and gratuitous ARP support.

To configure the pseudowire subscriber interface device:

1. Specify that you want to configure the pseudowire subscriber logical interface device.

```
user@host# edit interfaces ps0
```

2. Specify the logical tunnel interface that is the anchor point for the pseudowire logical interface device. The anchor point must be an **lt** device in the format **lt-fpc/pic/port**.



CAUTION: Do not reconfigure the logical tunnel interface that is associated with the pseudowire subscriber interface device unless you first deactivate all subscribers that are using the pseudowire subscriber interface.



NOTE: Tunnel services must be enabled on the **lt** interface that is the anchor point. You use the command, **set chassis fpc slot-number pic pic-number tunnel-services bandwidth bandwidth** to enable tunnel services.

```
[edit interfaces ps0]
```

```
user@host# set anchor-point lt-1/0/10
```

3. (Optional) Specify the MAC address for the pseudowire logical interface device.



NOTE: You should ensure that you change the MAC address prior to passing traffic or binding subscribers on the pseudowire port. Changing the MAC address when the pseudowire port is active (for example, while an upper layer protocol is negotiating) can negatively impact network performance until adjacencies learn of the new MAC address.

```
[edit interfaces ps0]
```

```
user@host# set mac 00:00:5E:00:53:55
```

4. (Optional) Specify the VLAN tagging method used for the pseudowire logical interface device. You can specify single tagging, dual (stacked) tagging, mixed (flexible) tagging, or no tagging.

```
[edit interfaces ps0]
```

```
user@host# set flexible-vlan-tagging
```

See *Enabling VLAN Tagging* for additional information about VLAN tagging.

5. (Optional) Specify the MTU for the pseudowire logical interface device. If you do not explicitly configure the MTU, the router uses the default value of 1500.

```
[edit interfaces ps0]
```

```
user@host# set mtu 2500
```

See *Setting the Protocol MTU* for additional information.

6. (Optional) Specify that the pseudowire logical interface device does not respond to gratuitous ARP requests.

```
[edit interfaces ps0]
user@host# set no-gratuitous-arp-request
```

See *Configuring Gratuitous ARP* for additional information.

7. (Optional) Specify that reverse-path forwarding checks are performed for traffic on the pseudowire logical interface device.

```
[edit interfaces ps0]
user@host# set rpf-check
```

See *Configuring Unicast RPF* for additional information.

8. Configure additional optional parameters for the pseudowire logical interface device, such as *description*, *apply-groups*, *apply-groups-except*, and *traceoptions*.



NOTE: You cannot dynamically change an anchor point that has active pseudowire devices stacked above it. If you need to change such an anchor point, you must perform the following steps:

1. Deactivate the stacked pseudowires and commit. This may require bringing down any subscribers using the pseudowires.

```
[edit interfaces]
user@host# deactivate psnumber
user@host# commit
```

2. Change the anchor on the deactivated pseudowire and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point lt-new-lt-interface-number
user@host# commit
```

3. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [Tunnel Interface Configuration on MX Series Routers Overview](#)
- [Router Chassis Configuration Statements](#)

Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface

This topic describes how to configure a pseudowire transport logical interface. A pseudowire device can have only one transport logical interface.

A pseudowire logical device and its related pseudowire logical interfaces are dependent on the state of the underlying logical transport interface device, which is either the Layer 2 VPN or Layer 2 circuit.



NOTE: We recommend that you use unit 0 to represent the transport logical interface for the pseudowire device. Non-zero unit numbers represent *service* logical interfaces used for pseudowire subscriber interfaces.

To configure a pseudowire transport logical interface:

1. Specify that you want to configure the pseudowire subscriber logical interface device.

```
[edit]
user@host# edit interfaces ps0
```

2. Specify that you want to configure unit 0, which represents the transport logical interface.

```
[edit interfaces ps0]
user@host# edit unit 0
```

3. Specify the **ethernet-ccc** encapsulation method for the transport logical interface.

```
[edit interfaces ps0 unit 0]
user@host# set encapsulation ethernet-ccc
```

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces

This topic describes the steps for configuring Layer 2 circuit signaling used for the pseudowire subscriber logical interface support. You can also use Layer 2 VPN signaling for pseudowire subscriber logical interfaces. The two methods are mutually exclusive; you can use only one method for a particular pseudowire.

To configure Layer 2 circuit signaling for pseudowire interfaces:

1. Specify that you want to configure Layer 2 circuit parameters at the protocols hierarchy level.

```
[edit protocols]
```

```
user@host# edit l2circuit
```

2. Specify the IP address of the neighbor, to identify the PE router used for the Layer 2 circuit.

```
[edit protocols l2circuit]
user@host# edit neighbor 192.168.102.15
```

3. Specify the interface used by the Layer 2 circuit traffic.

```
[edit protocols l2circuit neighbor 192.168.102.15]
user@host# edit interface ps1.0
```

4. Configure the virtual circuit ID that identifies the Layer 2 circuit for the pseudowire.

```
[edit protocols l2circuit neighbor 192.168.102.15 interface ps1.0]
user@host# set virtual-circuit-id 5
```

5. (Optional) If multiple VLAN interfaces are carried over the pseudowire Layer 2 payload, configure the **no-vlan-id-validate** statement. This statement prevents VLAN validation in the signaling.

```
[edit protocols l2circuit neighbor 192.168.102.15 interface ps1.0]
user@host# set no-vlan-id-validate
```

For more information about Layer 2 circuits, see *Configuring Interfaces for Layer 2 Circuits*.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [Configuring Interfaces for Layer 2 Circuits](#)

Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces

This topic describes the steps for configuring Layer 2 VPN signaling used for the pseudowire subscriber logical interface support. You can also use Layer 2 circuit signaling for pseudowire subscriber logical interfaces. The two methods are mutually exclusive; you can use only one method on a particular pseudowire.

To configure Layer 2 VPN signaling for pseudowire interfaces:

1. Specify the name of the routing instance you want to configure.

```
[edit]
user@host# edit routing-instances l2vpn0
```

2. Configure the Layer 2 VPN routing instance type.

```
[edit routing-instances l2vpn0]
user@host# set instance-type l2vpn
```

3. Associate the pseudowire logical interface for the Layer 2 VPN.

```
[edit routing-instances l2vpn0]  
user@host# set interface ps1.0
```

4. Configure the unique identifier for the routes that belong to the Layer 2 VPN.

```
[edit routing-instances l2vpn0]  
user@host# set route-distinguisher 198.51.100.101100
```

5. Configure the VPN routing and forwarding (VRF) target of the routing instance.

```
[edit routing-instances l2vpn0]  
user@host# set vrf-target target:10:100
```

6. Specify that you want to configure the Layer 2 VPN protocol for the routing instance.

```
[edit routing-instances l2vpn0]  
user@host# edit protocols l2vpn
```

7. Configure the encapsulation type for the routing instance.

```
[edit routing-instances l2vpn0 protocols l2vpn]  
user@host# set encapsulation-type ethernet
```

8. Specify the site name and site identifier for the Layer 2 VPN.

```
[edit routing-instances l2vpn0 protocols l2vpn]  
user@host# set site PE1 site-identifier 1
```

9. Specify the interface that connects to the site, and the remote interface to which you want the specified interface to connect.

```
[edit routing-instances l2vpn0 protocols l2vpn]  
user@host# set interface ps1.0 remote-site-id 2
```

10. Configure the tracing options for traffic that uses the Layer 2 VPN.

```
[edit routing-instances l2vpn0 protocols l2vpn]  
user@host# set traceoptions file l2vpn flag all
```

**Related
Documentation**

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface

This topic describes how to configure a pseudowire service logical interface. Service logical interfaces represent the attachment circuits for pseudowire logical interfaces.

As described in the “[Pseudowire Subscriber Logical Interfaces Overview](#)” on page 307, you can choose whether to configure a service logical interface together with a higher subscriber logical interface, depending upon the business need. In a broadband edge configuration, the higher subscriber logical interface is the demarcation point for subscribers. However, in a business edge configuration, the service logical interface is the demarcation point for the business subscribers, and also serves as the subscriber logical interface, so no subscriber logical interfaces are explicitly configured.



NOTE: Non-zero unit numbers represent *service* logical interfaces used for pseudowire subscriber interfaces. Use unit 0 to represent the *transport* logical interface for the pseudowire device.

To configure a pseudowire service logical interface:

1. Specify that you want to configure the pseudowire subscriber logical interface device.

```
[edit]
user@host# edit interfaces ps0
```

2. Configure the unit for the service logical interface. Use a non-zero unit number.

```
[edit interfaces ps0]
user@host# edit unit 1
```

3. Configure the VLAN tag IDs.

```
[edit interfaces ps0 unit 1]
user@host# set vlan-tags outer 1 inner 1
```

4. Configure the interface to respond to ARP requests when the device has an active route to the ARP request target address.

```
[edit interfaces ps0 unit 1]
user@host# set proxy-arp
```

5. Specify that you want to configure the protocol family information. Pseudowire service logical interfaces support IPv4 (inet), IPv6 (inet6), and PPPoE (pppoe) protocol families.

For example, to configure the IPv4 family:

- a. Specify that you want to configure IPv4.

```
[edit interfaces ps0 unit 1]
user@host# edit family inet
```

b. Configure the parameters for the family.

```
[edit interfaces ps0 unit 1 family inet]
user@host# set filter input filter 1 output filter 4
user@host# set mac-validate loose
user@host# set input-hierarchical-policer policer-1
user@host# set unnumbered-address lo0.0 preferred-source-address 198.51.100.11
```

**Related
Documentation**

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)

CHAPTER 28

Configuring Hierarchical CoS Pseudowire Subscriber Interfaces

- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)

Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

Junos OS supports two aspects of CoS for MPLS pseudowire subscriber interfaces. You can apply CoS rewrite rules and behavior aggregate (BA) classifiers to MPLS pseudowire subscriber interfaces. In addition, CoS performs egress hierarchical shaping towards the subscriber on MPLS pseudowire subscriber interfaces.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS is supported on MX Series routers with either EQ DPCs or MPC/MICs installed.

On Juniper Networks MX Series routers, MPC/MIC and EQ DPC interfaces support a four-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), the interface set or the underlying interface (level 2), one or more logical interfaces (level 3), and one or more queues (level 4). Although all CoS scheduling hierarchies are four-level, level 1 is always the physical interface and level 4 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure; for example, a logical interface or an interface set and where those interfaces reside in the scheduling hierarchy, either level 2 or level 3. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling* and *three-level hierarchical scheduling* in this discussion.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [*Understanding Hierarchical CoS for Subscriber Interfaces*](#)
- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 327](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)

- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [hierarchical-scheduler on page 493](#)

CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces

CoS supports two-level and three-level hierarchies for MPLS pseudowire subscriber interfaces.

To configure two-level scheduling, include the **maximum-hierarchy-levels 2** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point.

To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point. Use the following guidelines for configuring the **implicit-hierarchy** option:

- If an output traffic-control profile is configured on the pseudowire transport interface and on a pseudowire service interface, the two interfaces form a scheduling hierarchy. The pseudowire transport interface resides in a level 2 scheduler node and the pseudowire service interface resides in a level 3 scheduler node.
- If an output traffic-control profile is configured on the pseudowire services interface but not on a pseudowire transport interface, the pseudowire services interface resides in a level 3 scheduler node.
- If an output traffic-control profile is only configured on the pseudowire transport interface and not on the pseudowire services interface, the pseudowire transport interface resides in a level 3 scheduler node and all pseudowire traffic uses this node.

If the **implicit-hierarchy** option is not set on the logical tunnel anchor point, logical interfaces behave normally with the hierarchical-scheduler mode configured with or without the **hierarchical-scheduler maximum-hierarchy-levels** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement. In this case, when you apply a traffic-control profile to the pseudowire and service logical interfaces, they both reside in level 3 scheduler nodes and do not form a scheduling hierarchy, which might not be the desirable behavior. In business edge, where only the pseudowire logical interfaces need to be shaped, applying the traffic-control profile at just the transport logical interface may be sufficient.

When configuring the logical tunnel physical interface for the maximum hierarchy level, all pseudowire logical interfaces operating on the physical interface use the same hierarchy model. If you want to mix two-level and three-level scheduling hierarchies, you can group the pseudowires together by hierarchy levels and share the same logical tunnel anchor point or you can use three-level scheduling for all pseudowires over the anchor point.

To specify rewrite rules and classifiers on pseudowire interfaces, reference the pseudowire device under the **[edit class-of-service interfaces]** hierarchy level and specify the rewrite rules and classifiers for the pseudowire interfaces.

To control all pseudowire traffic using the same logical tunnel interface, apply CoS policies at the physical interface for the anchor logical tunnel.



NOTE: Starting with Junos OS Release 17.3R1, stateful anchor point redundancy support is provided for pseudowire subscriber logical interface by the underlying redundant logical tunnel interface (rlt) in active-backup mode. This redundancy protects the access and the core facing link against anchor PFE (Packet Forwarding Engine) failure. Starting in Junos OS Release 18.1R1, full hierarchical CoS support is provided for stateful anchor point redundancy of pseudowire subscriber logical interfaces. Both transport and services logical interfaces created for the pseudowire subscriber logical interface are stacked on the underlying redundant logical tunnel control logical interface. This logical interface stacking model is used for both redundant and non-redundant pseudowire subscriber logical interfaces. Hierarchical CoS is supported and configured the same on both redundant and non-redundant pseudowire subscriber logical interfaces.

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, full hierarchical CoS support is provided for stateful anchor point redundancy of pseudowire subscriber logical interfaces.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [*Understanding Hierarchical CoS for Subscriber Interfaces*](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 329](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 334](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 336](#)
- [Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview on page 310](#)
- [hierarchical-scheduler on page 493](#)

Configuring CoS Two-Level Hierarchical Scheduling

- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 327](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 329](#)

CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two. In a two-level scheduling hierarchy, all logical interfaces and interface sets share a single level 2 node. [Table 26 on page 327](#) summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.


Table 26: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes

Level 1	Level 2	Level 3	Level 4
Physical interface	—	Pseudowire transport logical interface	One or more queues
Physical interface	—	Interface set	One or more queues
Physical interface	—	Pseudowire service logical interface	One or more queues

You use the two-level hierarchical scheduling when you have many pseudowires but you do not require shaping specific to the subscriber logical interface. For example, when your configuration is one subscriber per pseudowire interface.

[Figure 17 on page 328](#) shows a two-level hierarchical scheduling configuration for the MPLS pseudowires. In this configuration, level 1 is the physical interface used for the logical tunnel anchor node. All of the pseudowire transport interfaces share a single level 2 node. The level 3 nodes are the pseudowire transport logical interfaces (ps0.0, ps1.0, and ps2.0). In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles.

Figure 17: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration



Two-level hierarchical scheduling has up to eight class of service queues. For this configuration, include the `maximum-hierarchy-levels 2` option under the `[edit interfaces interface-name hierarchical-scheduler]` hierarchy level at the physical interface for the anchor logical tunnel.



NOTE: You cannot configure shaping policies on both the pseudowire logical interfaces and the subscriber logical interfaces over the same pseudowire. If a traffic-control profile is configured on a pseudowire logical interface, and CoS policies are configured on the subscriber logical interface over another pseudowire, all of the logical interfaces are at level 3 and act as peers.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 329](#)
- [hierarchical-scheduler on page 493](#)

Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces

Before configuring CoS parameters for MPLS pseudowire subscriber interfaces, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See [“Configuring a Pseudowire Subscriber Logical Interface” on page 313](#).
2. Configure the pseudowire device count. See [“Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router” on page 315](#).
3. Configure the pseudowire device including the logical tunnel anchor point. See [“Configuring a Pseudowire Subscriber Logical Interface Device” on page 315](#).
4. Configure the pseudowire transport logical interface. See [“Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface” on page 318](#).
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See [“Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces” on page 318](#) or [“Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces” on page 319](#).
6. Configure the pseudowire logical interfaces. See [“Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface” on page 321](#).

To configure CoS policies on MPLS pseudowire subscriber interfaces using two-level scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For two-level scheduling the hierarchical scheduler must be set to **maximum-scheduler levels 2**.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler maximum-hierarchy-levels 2
```

2. Specify the traffic-control profile to use on the pseudowire logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

4. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

5. Apply the rewrite rule and classifier to the pseudowire interface.

For the *interface_name* parameter, specify the pseudowire device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

**Related
Documentation**

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- *Mapping CoS Component Inputs to Outputs*
- *Understanding Hierarchical CoS for Subscriber Interfaces*
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 327](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [hierarchical-scheduler on page 493](#)

Configuring CoS Three-Level Hierarchical Scheduling

- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 334](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 336](#)

CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a scheduling hierarchy. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set and the use of underlying interfaces. In all variations, the physical interface on which the logical tunnel resides is a level 1 CoS scheduler node and the queues reside at level 4. Three-level scheduling hierarchies can have up to eight class of service queues.

[Table 27 on page 331](#) summarizes the most common three-level hierarchical scheduling configurations and shows the interface hierarchy and CoS scheduler nodes.

Table 27: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels

Level 1	Level 2	Level 3	Level 4
Physical interface	Pseudowire interface set	Pseudowire service logical interfaces	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire interface set	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire service logical interfaces	One or more queues


Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface

Figure 18 on page 332 shows an MPLS pseudowire three-level scheduling hierarchy that includes two pseudowire service logical interfaces over a pseudowire transport logical interface. This variation uses the following scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2) for subscriber sessions
- Level 2—Pseudowire transport logical interface (ps0.0)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles at the pseudowire transport logical interfaces (level 2) and the pseudowire service logical interfaces (level 3).

Figure 18: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface




Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set

Figure 19 on page 333 shows another variation of MPLS pseudowire three-level hierarchical scheduling that includes two pseudowire service logical interfaces over a pseudowire service interface set. This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2)
- Level 2—Pseudowire service interface set
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profile at the pseudowire service interfaces (level 3) and at the interface set (level 2). This variation is most useful for subscriber edge deployments.

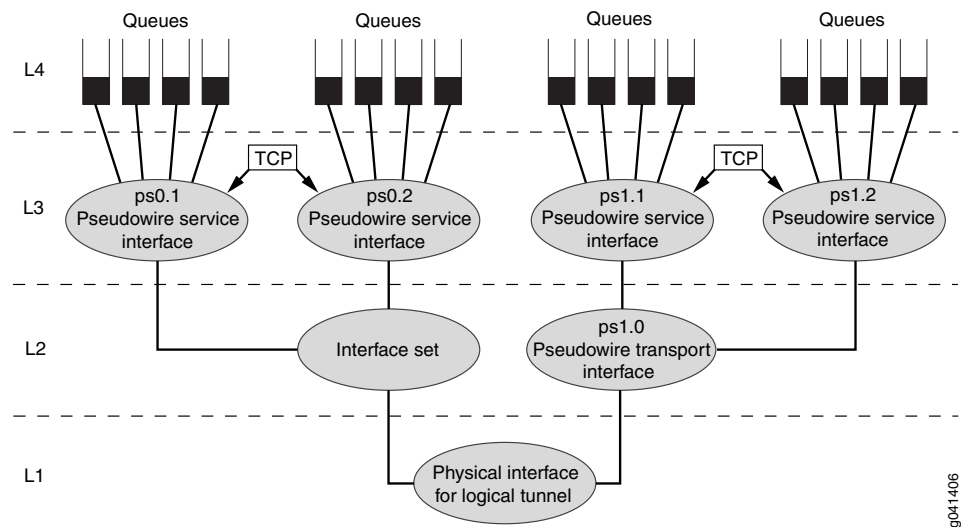
Figure 19: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set



Three-Level Scheduling Hierarchy Combined Deployment Scenario

Figure 20 on page 333 shows a deployment scenario that combines the three-level hierarchical scheduling scenarios in Figure 18 on page 332 and Figure 19 on page 333.

Figure 20: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces—Deployment Scenario



This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1, ps0.2, ps1.1, and ps1.2)

- Level 2—Service interface set for pseudowire service interfaces (ps0.1 and ps0.2) and transport logical interface (ps1.0) for the pseudowire service logical interfaces (ps1.1 and ps1.2)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles to the interfaces at both level 2 and level 3, as well as the interface set at level 2.

Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
- [*Understanding Hierarchical CoS for Subscriber Interfaces*](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 334](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 336](#)
- [hierarchical-scheduler on page 493](#)

Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface)

Before configuring CoS three-level scheduling on pseudowire logical interfaces over a transport logical interface, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See [“Configuring a Pseudowire Subscriber Logical Interface” on page 313](#).
2. Configure the pseudowire device count. See [“Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router” on page 315](#).
3. Configure the pseudowire device including the logical tunnel anchor point. See [“Configuring a Pseudowire Subscriber Logical Interface Device” on page 315](#).
4. Configure the pseudowire transport logical interface. See [“Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface” on page 318](#).
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See [“Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces” on page 318](#) or [“Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces” on page 319](#).
6. Configure the pseudowire logical interfaces. See [“Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface” on page 321](#).

Three-level scheduling on pseudowire logical interfaces over a transport logical interface requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire transport logical interface. To configure CoS policies on three-level scheduling on pseudowire logical interfaces over a transport logical interface:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler implicit-hierarchy
```

2. Specify the traffic-control profile to use on the pseudowire logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Specify the traffic-control profile to use on the pseudowire transport logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

4. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

5. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

6. Apply the rewrite rule and classifier to the pseudowire interfaces.

For the *interface_name* parameter, specify the pseudowire device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
```

```
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
```

```
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

**Related
Documentation**

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- *Mapping CoS Component Inputs to Outputs*
- *Understanding Hierarchical CoS for Subscriber Interfaces*
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 336](#)
- [hierarchical-scheduler on page 493](#)

Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set)

Before configuring three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set, you must first complete the following tasks:

1. Configure the pseudowire logical interfaces. See [“Configuring a Pseudowire Subscriber Logical Interface” on page 313](#).
2. Configure the pseudowire device count. See [“Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router” on page 315](#).
3. Configure the pseudowire device including the logical tunnel anchor point. See [“Configuring a Pseudowire Subscriber Logical Interface Device” on page 315](#).
4. Configure the pseudowire transport logical interface. See [“Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface” on page 318](#).
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See [“Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces” on page 318](#) or [“Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces” on page 319](#).
6. Configure the pseudowire logical interfaces. See [“Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface” on page 321](#).

Three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire logical interface-set. To configure CoS policies on MPLS pseudowire subscriber interfaces using three-level implicit hierarchical scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler implicit-hierarchy
```

2. Specify the traffic-control profile to use on the pseudowire logical interfaces.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Define a pseudowire logical interface set and configure the traffic-control profile used for the interface set.

```
[edit class-of-service]
user@host#edit interfaces
user@host#edit interface-set interface-set-name
user@host#edit output-traffic-control-profile profile-name
```

4. Group the pseudowire logical interfaces in the pseudowire logical interface set.

```
[edit ]
user@host#edit interfaces
user@host#edit interface-set interface-set-name
user@host#edit interface ps ps-device-name
user@host#edit unit logical-unit-number
```

5. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

6. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
```

```
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

7. Apply the rewrite rule and classifier to the pseudowire interfaces.

For the *interface_name* parameter, specify the ps device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

**Related
Documentation**

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- *Mapping CoS Component Inputs to Outputs*
- *Understanding Hierarchical CoS for Subscriber Interfaces*
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 331](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 324](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 334](#)
- [hierarchical-scheduler on page 493](#)

PART 6

Configuring Wi-Fi Access Gateway

- [Configuring Wi-Fi Access Gateway on page 341](#)

CHAPTER 31

Configuring Wi-Fi Access Gateway


- [Wi-Fi Access Gateway Overview on page 342](#)
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Wi-Fi Access Gateway Configuration Overview on page 346](#)
- [Configuring a Psuedowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway on page 347](#)
- [Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation on page 349](#)
- [Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways on page 350](#)
- [Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways on page 354](#)

Wi-Fi Access Gateway Overview

Wi-Fi access gateway (WAG) provides the public with Wi-Fi access from a residential Wi-Fi network or from a business Wi-Fi network. At home, subscribers have their existing Wi-Fi network; however, a part of their network is available for the general public to use. Members of the public who have an account with the same Internet service provider as the subscriber has at home can access the Internet and mobile network through the public part of the subscriber's Wi-Fi connection when they are in close proximity to the subscriber's home. WAG authenticates and connects subscribers regardless of their physical location.

Service providers can deploy the MX Series router as a broadband network gateway (BNG) within their network, and then deploy the BNG as a WAG as shown in [Figure 21 on page 342](#).

Figure 21: MX Series Router Deployed as a WAG



After a WAG has been deployed, service providers can configure the WAG to create secure wireless home network connections for computers, laptops, and other Wi-Fi electronic products (such as game systems, tablets, or mobile phones). WAG offers wireline and mobile service providers the following deployments and business value opportunities:

- **Wireline service providers**—The WAG deployment is based on an in-home division of access points or public access points, and works with any Wi-Fi access point that creates a generic routing encapsulation (GRE) tunnel to the MX Series router. This deployment protects subscribers and reduces churn by including free Wi-Fi with a paid wireline subscription. For added value, service providers can also sell ad hoc access or mode, such as airport, public safety, search-and-rescue, and café access.
- **Mobile service providers**—The WAG deployment is based on the mobile service provider's own access points, or wholesale and retail with the wireline service provider. Service providers that offer quadruple play, where TV, Internet, wireless, and landline phone services are combined, can leverage both wireline and wireless assets. This deployment offsets costs in mobile packet core and radio access network infrastructures with the

ability to offload mobile data. For added value, service providers can offer Wi-Fi for all devices with a mobile data plan as a competitive differentiator.

Customers who purchase broadband can also receive Wi-Fi on any community Wi-Fi access point. Subscribers have a private and secure home connection, and can also access a public connection that is shared by other subscribers. To maintain a level of security and protect the private home connection, the two networks are separated. This separation ensures a strong level of bandwidth on the subscribers' personal connections.

Subscriber services such as authentication, authorization, and accounting (AAA); address assignment; hierarchical quality of service (QoS); lawful intercept; and class of service (CoS) are supported for individual Dynamic Host Configuration Protocol (DHCP) subscribers within the GRE tunnels. Using GRE tunnels for Wi-Fi provides the following benefits:


- Wi-Fi users who are not directly connected through Layer 2 to WAG are authenticated because GRE tunnels transmit Layer 2 information across any IP network.
- Services based on user equipment-specific information are applied using the media access control (MAC) address or Subscriber Identity Module (SIM) card.
- Services are applied in the network, not just at the Wi-Fi access point.
- The soft GRE or Ethernet-over-GRE standard is supported on most Wi-Fi access points. For services using the Ethernet over GRE standard, only one side of the tunnel needs to be configured; the other end learns the remote IP addresses of all remote tunnel endpoints by examining the incoming GRE packets.

- Related Documentation**
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
 - [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Wi-Fi Access Gateway Deployment Model Overview

Figure 22 on page 344 shows an MX Series router broadband network gateway (BNG) deployed as a Wi-Fi access gateway (WAG). The WAG provides a multiservice edge with a full broadband feature set that is highly reliable because of the included redundant hardware. Ethernet frames from the user equipment device must be tunneled to the BNG across an IP cloud or public Internet.

Figure 22: MX Series as Wi-Fi Access Gateway Deployment Model



To support the MX Series BNG deployed as a WAG, dynamic-bridged generic routing encapsulation (GRE) tunnels are created and terminated at the BNG when it receives GRE traffic from the wireless access point (WAP). Dynamic Host Configuration Protocol (DHCP) subscribers are transported through GRE tunnels as either VLAN-tagged per service set identifier (SSID) or untagged. When the user equipment device connects to the SSID and begins to send traffic, the access point initiates a Layer 2 soft GRE or Ethernet-over-GRE connection to the MX Series BNG and the BNG dynamically builds the GRE tunnel. GRE tunnels are cleared after all of the subscribers within a GRE tunnel have logged out and a configurable timer has expired.

This deployment model supports a full set of services per user equipment device and per access point. Subscriber services such as authentication, authorization, and accounting (AAA); address assignment; hierarchical quality of service (QoS); lawful intercept; and class of service (CoS) are supported for individual DHCP subscribers within the GRE tunnels. No additional service cards are required for GRE or QoS because all features run inline on MPCs.

External RADIUS proxy supports Extensible Authentication Protocol (EAP) Subscriber Identity Module (SIM), Tunneled Transport Layer Security (TTLS), and Authentication and Key Agreement (AKA) protocols. The External RADIUS proxy also integrates with HTTP redirect to the Web portal.

The MX Series as WAG deployment model also supports the wholesale of access point access to multiple retail service providers. This wholesaling allows the local breakout of traffic or Layer 3 handoff to retail service providers.

Related Documentation

- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway

Dynamic-bridged generic routing encapsulation (GRE) tunnels and the Wi-Fi access gateway support interface stacks for VLAN-tagged and untagged subscribers. Subscriber features such as dynamic and service profiles for DHCP subscribers, lawful intercept, firewall filters, and change of authorization (CoA) are supported.

Scaling limitations of pseudowire subscriber interface devices (*psn* IFDs) require that multiple tunnels share the same *psn* IFD. The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD).



NOTE: The *psn* IFD used to service dynamic GRE tunnel terminations cannot be simultaneously used to service MPLS pseudowire terminations.

Subscriber services and lawful intercept are supported only at the IP demultiplexing (demux) interface level.



NOTE: A GRE tunnel cannot have both untagged and tagged subscribers.

The tagged model and the untagged model are described in the following sections:

- [Dynamic VLAN-Tagged Subscribers on page 345](#)
- [Untagged Subscribers on page 345](#)

Dynamic VLAN-Tagged Subscribers

To make provisioning and troubleshooting easier for VLAN-tagged subscribers, use the same set of VLANs on all of the Wi-Fi access points. Doing this requires that the same pseudowire subscriber interface service logical interface (*psn* IFL) (associated with a VLAN ID) on a *psn* IFD represents multiple GRE tunnels.

A dynamic VLAN demux interface (demux0.yyyyyyyy) is created for each VLAN tag and is stacked over the tunnel *psn* interface (*psn.xxxxxxxx*). There can be multiple VLANs (single and dual-tagged) over the same GRE tunnel. The subscribers' IP demux interfaces are then created over the VLAN demux interface.

Untagged Subscribers

Untagged DHCP subscribers can be created directly over the GRE tunnel. For each subscriber, an IP demux interface (demux0.yyyyyyyy) is created and is stacked over the

tunnel *psn* logical interface (*psn.xxxxxxxx*). There can be multiple subscribers over the same GRE tunnel.

- Related Documentation**
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
 - [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Wi-Fi Access Gateway Configuration Overview

To configure the MX Series router as a Wi-Fi access gateway (WAG):

1. Configure a pseudowire subscriber logical interface device.

See “Configuring a Psuedowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway” on page 347.

2. Configure the conditions for enabling dynamic-bridged GRE tunnels.

See “Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation” on page 349.

3. Configure the type of dynamic-bridged GRE tunnel that carries subscriber traffic to the WAG:



NOTE: A GRE tunnel cannot have both untagged and tagged subscribers.

- If the subscriber traffic is VLAN-tagged, see “Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways” on page 350.
- If the subscriber traffic is untagged, see “Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways” on page 354.

- Related Documentation**
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
 - [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)

Configuring a Pseudowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway

Before you begin, you must create a logical tunnel interface:

- Configure the maximum number of pseudowire logical interfaces devices. See [“Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router” on page 315](#).
- Configure a tunnel interface. See *Tunnel Interface Configuration on MX Series Routers Overview*.

To configure the pseudowire subscriber logical interface device on which the dynamic-bridged GRE tunnel is built on the MX Series router Wi-Fi access gateway:

1. Specify that you want to configure the pseudowire subscriber logical interface device.

```
user@host# edit interfaces psn
```

For example:

```
user@host# edit interfaces ps0
```

2. Specify the logical tunnel interface that is the anchor point for the pseudowire logical device interface.

```
[edit interfaces psn]
user@host# set anchor-point lt-fpc/pic/port
```

For example:

```
[edit interfaces ps0]
user@host# set anchor-point lt-0/0/0
```

3. Configure three-level hierarchical scheduling on the logical tunnel interface.

```
[edit interfaces lt-fpc/pic/port]
user@host# set hierarchical-scheduler implicit-hierarchy
```

For example:

```
[edit interfaces lt-0/0/0]
user@host# set hierarchical-scheduler implicit-hierarchy
```

4. Configure the mixed VLAN tagging method for the pseudowire logical interface device.

```
[edit interfaces psn]
user@host# set flexible-vlan-tagging
```



NOTE: You must configure `flexible-vlan-tagging` even if only untagged subscriber packets are being transported on the dynamic-bridged GRE tunnel.

For example:

```
[edit interfaces ps0]  
user@host# set flexible-vlan-tagging
```

5. Specify that you want to configure unit 0, which represents the transport logical interface.

```
[edit interfaces psn]  
user@host# edit unit 0
```

For example:

```
[edit interfaces ps0]  
user@host# edit unit 0
```

6. Specify the Ethernet CCC encapsulation method for the transport logical interface.

```
[edit interfaces psn unit 0]  
user@host# set encapsulation ethernet-ccc
```

For example:

```
[edit interfaces ps0 unit 0]  
user@host# set encapsulation ethernet-ccc
```

**Related
Documentation**

- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation

Before you begin:

- Configure the pseudowire logical device on which to build the dynamic-bridged GRE tunnel. See [“Configuring a Pseudowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway” on page 347](#).
- Configure interface lo0 with the source IP address of the GRE tunnels for the Wi-Fi access gateway (WAG). Use the IP address of the MX Series router that you want to receive the incoming GRE traffic. This address cannot be the primary or preferred address on lo0. See *Configuring a Loopback Interface*.

To configure the conditions for enabling dynamic-bridged generic routing encapsulation (GRE) tunnel creation on the MX Series router WAG, you configure one or more GRE tunnel groups. Multiple GRE tunnel groups can have the same **source-address** or the same **destination-networks** value, but you cannot use a specific **source-address** and **destination-networks** combination in more than one GRE tunnel group.

To configure a GRE tunnel group:

1. Name the dynamic GRE tunnel group.

```
[edit services]
user@host# set soft-gre group-name
```

For example:

```
[edit services]
user@host# set soft-gre AP-Group1
```

2. Specify the source IP address of the GRE tunnels for the WAG. Use the IP address of the MX Series router that you configured to receive the incoming GRE traffic.

```
[edit services soft-gre group-name]
user@host# set source-address wag-ip-address
```

For example:

```
[edit services soft-gre AP-Group1]
user@host# set source-address 192.168.0.20
```

3. Specify the IP subnets from which GRE traffic can be processed.

```
[edit services soft-gre group-name]
user@host# set destination-networks [prefix]
```

For example:

```
[edit services soft-gre AP-Group1]
user@host# set destination-networks 20.0.0.0/24
```

4. Specify the pseudowire subscriber interface device (IFD) on which to build the dynamic-bridged GRE tunnels.

```
[edit services soft-gre group-name]
user@host# set service-interface psn
```

For example:

```
[edit services soft-gre AP-Group1]
user@host# set service-interface ps0
```

5. Specify the dynamic profile that configures the GRE tunnel.

```
[edit services soft-gre group-name]
user@host# set dynamic-profile profile-name
```

For example:

```
[edit services soft-gre AP-Group1]
user@host# set dynamic-profile tunnel_profile
```

6. (Optional) Configure the number of seconds that a GRE tunnel remains up after the last subscriber session on the tunnel has ended.

```
[edit services soft-gre group-name]
user@host# set tunnel-idle-timeout seconds
```

The default **tunnel-idle-timeout** value is 120 seconds.

For example:

```
[edit services soft-gre AP-Group1]
user@host# set tunnel-idle-timeout 60
```

7. To configure another GRE tunnel group, repeat this procedure.

Related Documentation

- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways

To configure subscriber interfaces for VLAN-tagged Dynamic Host Configuration Protocol (DHCP) subscribers on dynamic-bridged generic routing encapsulation (GRE) tunnels:

1. Name the dynamic profile.

```
[edit]
user@host# set dynamic-profiles profile-name
```

For example:

```
[edit]
user@host# set dynamic-profiles tunnel_profile
```

2. Define the interface with the internal variable used by the router to match the interface name of the receiving interface.

```
[edit dynamic-profiles profile-name]
user@host# edit interfaces $junos-interface-ifd-name
```

For example:

```
[edit dynamic-profiles tunnel_profile]
user@host# edit interfaces $junos-interface-ifd-name
```

3. Define the unit with the internal variable.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name]
user@host# set unit $junos-interface-unit
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name]
user@host# set unit $junos-interface-unit
```

4. Define the unit family type.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set family (inet | inet6)
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set family inet
```

5. Enable the local address for the interface to be derived from the loopback interface address.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit family (inet | inet6)]
user@host# set unnumbered-address lo0.0
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit family inet]
user@host# set unnumbered-address lo0.0
```

6. Configure the router to respond to any ARP request.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set proxy-arp
```

7. Configure stacked VLAN processing:

- a. Access the VLAN range configuration for stacked VLANs.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
```

```
user@host# edit auto-configure stacked-vlan-ranges
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# edit auto-configure stacked-vlan-ranges
```

- b. Specify the dynamic profile that is used to create VLANs.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges]
user@host# edit dynamic-profile profile-name
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges]
user@host# edit dynamic-profile auto_svlan_demux
```

- c. Specify that the VLAN dynamic profile accepts any type of VLAN Ethernet packet.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges dynamic-profile
profile-name]
user@host# set accept any
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges dynamic-profile
auto_svlan_demux]
user@host# set accept any
```

- d. Specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges dynamic-profile
profile-name]
user@host# set ranges low-tag-high-tag,low-tag-high-tag
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure stacked-vlan-ranges dynamic-profile
auto_svlan_demux]
user@host# set ranges 1000-1100,1200-1300
```

8. Configure single-tagged VLAN processing:

- a. Access the VLAN range configuration for single VLANs.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# edit auto-configure vlan-ranges
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# edit auto-configure vlan-ranges
```

- b. Specify the dynamic profile used to create VLANs.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges]
user@host# edit dynamic-profile profile-name
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges]
user@host# edit dynamic-profile auto_vlan_demux
```

- c. Specify that the VLAN dynamic profile accepts any type of VLAN Ethernet packet.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges dynamic-profile profile-name]
user@host# set accept any
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges dynamic-profile
auto_vlan_demux]
user@host# set accept any
```

- d. Specify the VLAN range that you want the dynamic profile to use.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges dynamic-profile profile-name]
user@host# set ranges low-tag-high-tag
```

For example:

```
[edit dynamic-profiles tunnel_profile interfaces $junos-interface-ifd-name unit
$junos-interface-unit auto-configure vlan-ranges dynamic-profile
auto_vlan_demux]
user@host# set ranges 1-50
user@host# set ranges 101-150
```

Related Documentation

- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Wi-Fi Access Gateway Configuration Overview on page 346](#)

Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways

To configure subscriber interfaces for untagged Dynamic Host Configuration Protocol (DHCP) subscribers on dynamic-bridged generic routing encapsulation (GRE) tunnels:

1. Name the dynamic profile.

```
[edit]
user@host# set dynamic-profiles profile-name
```

For example:

```
[edit]
user@host# set dynamic-profiles tunnel_demux
```

2. Define the interface with the internal variable used by the router to match the interface name of the receiving interface.

```
[edit dynamic-profiles profile-name]
user@host# edit interfaces $junos-interface-ifd-name
```

For example:

```
[edit dynamic-profiles tunnel_demux]
user@host# edit interfaces $junos-interface-ifd-name
```

3. Define the unit with the internal variable.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name]
user@host# set unit $junos-interface-unit
```

For example:

```
[edit dynamic-profiles tunnel_demux interfaces $junos-interface-ifd-name]
user@host# set unit $junos-interface-unit
```

4. Configure the variable for the underlying interface of the demux interfaces.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

For example:

```
[edit dynamic-profiles tunnel_demux interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

5. Define the unit family type.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set family (inet | inet6)
```

For example:

```
[edit dynamic-profiles tunnel_demux interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host# set family inet
```

- Related Documentation**
- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
 - [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
 - [Wi-Fi Access Gateway Configuration Overview on page 346](#)

PART 7

Troubleshooting

- [Configuring PPP Log Files on page 359](#)
- [Configuring PPP Trace Flags and Operations on page 363](#)
- [Configuring L2TP Log Files on page 367](#)
- [Configuring L2TP Trace Flags and Operations on page 371](#)
- [Contacting Juniper Networks Technical Support on page 375](#)

CHAPTER 32

Configuring PPP Log Files

- [Configuring the Number and Size of PPP Service Log Files on page 359](#)
- [Configuring Access to the PPP Service Log File on page 360](#)
- [Configuring the Severity Level to Filter Which PPP Service Messages Are Logged on page 360](#)
- [Configuring a Regular Expression for PPP Service Messages to Be Logged on page 361](#)

Configuring the Number and Size of PPP Service Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format *.number.gz*. The newest archived file is *.0.gz* and the oldest archived file is *.(maximum number)-1.gz*. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can 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 compressed and renamed *filename.0.gz*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until there are 20 trace files. Then the oldest file, *filename.19.gz*, is simply overwritten when the next oldest file, *filename.18.gz* is compressed and renamed to *filename.19.gz*.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

```
[edit protocols ppp-service traceoptions]  
user@host# set file ppp-service_1_logfile_1 files 20 size 2097152
```

Related Documentation

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring Access to the PPP Service Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

```
[edit protocols ppp-service traceoptions]  
user@host# set file ppp-service_1_logfile_1 world-readable
```

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

```
[edit protocols ppp-service traceoptions]  
user@host# set file ppp-service_1_logfile_1 no-world-readable
```

Related Documentation

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring the Severity Level to Filter Which PPP Service Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the

messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

```
[edit protocols ppp-service traceoptions]  
user@host# set level severity
```

**Related
Documentation**

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring a Regular Expression for PPP Service Messages to Be Logged

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

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

```
[edit protocols ppp-service traceoptions]  
user@host# set file ppp-service_1_logfile_1 match regex
```

**Related
Documentation**

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring PPP Trace Flags and Operations

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)
- [Configuring the PPP Service Trace Log Filename on page 364](#)
- [Configuring the PPP Service Tracing Flags on page 364](#)
- [Configuring Subscriber Filtering for PPP Service Trace Operations on page 365](#)

Tracing PPP Service Operations for Subscriber Access

The Junos OS trace feature tracks PPP service operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **jpppd**. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file **filename** reaches 128 kilobytes (KB), it is compressed and renamed **filename.0.gz**. Subsequent events are logged in a new file called **filename**, until it reaches capacity again. At this point, **filename.0.gz** is renamed **filename.1.gz** and **filename** is compressed and renamed **filename.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the [System Log Explorer](#).)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPP service tracing operations:

1. (Optional) Configure a trace log filename.
See [“Configuring the PPP Service Trace Log Filename” on page 364.](#)
2. (Optional) Configure the number and size of trace logs.
See [“Configuring the Number and Size of PPP Service Log Files” on page 359.](#)
3. (Optional) Configure user access to trace logs.
See [“Configuring Access to the PPP Service Log File” on page 360.](#)
4. (Optional) Configure a regular expression to filter the information to be included in the trace log.
See [“Configuring a Regular Expression for PPP Service Messages to Be Logged” on page 361.](#)
5. (Optional) Configure flags to specify which events are logged.
See [“Configuring the PPP Service Tracing Flags” on page 364.](#)
6. (Optional) Configure a severity level for messages to specify which event messages are logged.
See [“Configuring the Severity Level to Filter Which PPP Service Messages Are Logged” on page 360.](#)

Configuring the PPP Service Trace Log Filename

By default, the name of the file that records trace output for PPP service is **jpppd**. You can specify a different name with the **file** option.

To configure the filename for PPP service tracing operations:

- Specify the name of the file used for the trace output.

```
[edit protocols ppp-service traceoptions]  
user@host# set file ppp-service_logfile_1
```

Related Documentation

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring the PPP Service Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit protocols ppp-service traceoptions]
user@host# set flag flag
```

**Related
Documentation**

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring Subscriber Filtering for PPP Service Trace Operations

You can apply filters to the PPP service to limit tracing to particular subscribers or domains. Subscriber filtering simplifies troubleshooting in a scaled environment by enabling you to focus on a reduced set of trace results.

For subscriber usernames that have the expected form of *user@domain*, you can filter on the user, the domain, or both. You can use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms to match a greater number of subscribers.



NOTE: You cannot filter results using a wildcard in the middle of the user or domain terms. For example, the following uses of the wildcard are not supported: *tom*25@example.com*, *tom125@ex*.com*.

When you enable filtering by username, traces that have insufficient information to determine the username are automatically excluded.

To configure subscriber filtering:

- Specify the filter.

```
[edit protocols ppp-service traceoptions]
user@host# set filter user user@domain
```

Consider the following examples of using the wildcard for filtering:

- Filter results for the specific subscriber with the username, *tom@example.com*.

```
[edit protocols ppp-service traceoptions]
user@host# set filter user tom@example.com
```

- Filter results for all subscribers whose username begins with *tom*.

```
[edit protocols ppp-service traceoptions]
user@host# set filter user tom*
```

- Filter results for all subscribers whose username ends with *tom*.

```
[edit protocols ppp-service traceoptions]
user@host# set filter user *tom
```

- Filter results for subscribers with the username *tom* at all domains beginning with *ex*.

```
[edit protocols ppp-service traceoptions]
```

```
user@host# set filter user tom@ex*
```

- Filter results for all subscribers at all domains that end with ample.com.

```
[edit protocols ppp-service traceoptions]
```

```
user@host# set filter user *ample.com
```

- Filter results for all subscribers whose username begins with tom at domains that end with example.com.

```
[edit protocols ppp-service traceoptions]
```

```
user@host# set filter user tom*.*example.com
```

**Related
Documentation**

- [Tracing PPP Service Operations for Subscriber Access on page 363](#)

Configuring L2TP Log Files

- [Configuring the Number and Size of L2TP Log Files on page 367](#)
- [Configuring Access to the L2TP Log File on page 368](#)
- [Configuring a Regular Expression for L2TP Messages to Be Logged on page 368](#)
- [Configuring the Severity Level to Filter Which L2TP Messages Are Logged on page 368](#)

Configuring the Number and Size of L2TP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format *.number.gz*. The newest archived file is *.0.gz* and the oldest archived file is *.(maximum number)-1.gz*. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can 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 compressed and renamed *filename.0.gz*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until there are 20 trace files. Then the oldest file, *filename.19.gz*, is simply overwritten when the next oldest file, *filename.18.gz* is compressed and renamed to *filename.19.gz*.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

```
[edit services l2tp traceoptions]  
user@host# set file l2tp_1_logfile_1 files 20 size 2097152
```

Related Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring Access to the L2TP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

```
[edit services l2tp traceoptions]  
user@host# set file l2tp_1_logfile_1 world-readable
```

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

```
[edit services l2tp traceoptions]  
user@host# set file l2tp_1_logfile_1 no-world-readable
```

Related Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring a Regular Expression for L2TP Messages to Be Logged

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

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

```
[edit services l2tp traceoptions]  
user@host# set file l2tp_1_logfile_1 match regex
```

Related Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring the Severity Level to Filter Which L2TP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

```
[edit services l2tp traceoptions]  
user@host# set level severity
```

**Related
Documentation**

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring L2TP Trace Flags and Operations

- [Tracing L2TP Operations for Subscriber Access on page 371](#)
- [Configuring the L2TP Trace Log Filename on page 372](#)
- [Configuring the L2TP Tracing Flags on page 372](#)
- [Configuring Subscriber Filtering for L2TP Trace Operations on page 373](#)

Tracing L2TP Operations for Subscriber Access

The Junos OS trace feature tracks L2TP operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.



NOTE: This topic refers to tracing L2TP operations on MX Series routers. To trace L2TP operations on M Series routers, see *Tracing L2TP Operations*.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **jl2tpd**. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file **filename** reaches 128 kilobytes (KB), it is compressed and renamed **filename.0.gz**. Subsequent events are logged in a new file called **filename**, until it reaches capacity again. At this point, **filename.0.gz** is renamed **filename.1.gz** and **filename** is compressed and renamed **filename.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the [System Log Explorer](#).)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure L2TP tracing operations:

1. (Optional) Configure a trace log filename.
See [“Configuring the L2TP Trace Log Filename” on page 372](#).
2. (Optional) Configure the number and size of trace logs.
See [“Configuring the Number and Size of L2TP Log Files” on page 367](#).
3. (Optional) Configure user access to trace logs.
See [“Configuring Access to the L2TP Log File” on page 368](#).
4. (Optional) Configure a regular expression to filter the information to be included in the trace log.
See [“Configuring a Regular Expression for L2TP Messages to Be Logged” on page 368](#).
5. (Optional) Configure flags to specify which events are logged.
See [“Configuring the L2TP Tracing Flags” on page 372](#).
6. (Optional) Configure a severity level for messages to specify which event messages are logged.
See [“Configuring the Severity Level to Filter Which L2TP Messages Are Logged” on page 368](#).

Configuring the L2TP Trace Log Filename

By default, the name of the file that records trace output for L2TP is `jl2tpd`. You can specify a different name with the `file` option.

To configure the filename for L2TP tracing operations:

- Specify the name of the file used for the trace output.

```
[edit services l2tp traceoptions]  
user@host# set file l2tp_logfile_1
```

Related Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring the L2TP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit services l2tp traceoptions]
user@host# set flag flag
```

Related
Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

Configuring Subscriber Filtering for L2TP Trace Operations

Starting in Junos OS Release 14.1, you can apply filters to L2TP to limit tracing to particular subscribers or domains. Subscriber filtering simplifies troubleshooting in a scaled environment by enabling you to focus on a reduced set of trace results.

For subscriber usernames that have the expected form of **user@domain**, you can filter on the user, the domain, or both. You can use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms to match a greater number of subscribers.



NOTE: You cannot filter results using a wildcard in the middle of the user or domain terms. For example, the following uses of the wildcard are not supported: tom*25@example.com, tom125@ex*.com.

When you enable filtering by username, traces that have insufficient information to determine the username are automatically excluded.

To configure subscriber filtering:

- Specify the filter.

```
[edit services l2tp traceoptions]
user@host# set filter user user@domain
```



NOTE: This syntax is different than the syntax used to filter subscribers on M Series routers.

Consider the following examples of using the wildcard for filtering:

- Filter results for the specific subscriber with the username, tom@example.com.

```
[edit services l2tp traceoptions]
user@host# set filter user tom@example.com
```

- Filter results for all subscribers whose username begins with tom.

```
[edit services l2tp traceoptions]
user@host# set filter user tom*
```

- Filter results for all subscribers whose username ends with tom.

```
[edit services l2tp traceoptions]  
user@host# set filter user *tom
```

- Filter results for subscribers with the username tom at all domains beginning with ex.

```
[edit services l2tp traceoptions]  
user@host# set filter user tom@ex*
```

- Filter results for all subscribers at all domains that end with ample.com.

```
[edit services l2tp traceoptions]  
user@host# set filter user *ample.com
```

- Filter results for all subscribers whose username begins with tom at domains that end with example.com.

```
[edit services l2tp traceoptions]  
user@host# set filter user tom*.*example.com
```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, you can apply filters to L2TP to limit tracing to particular subscribers or domains. Subscriber filtering simplifies troubleshooting in a scaled environment by enabling you to focus on a reduced set of trace results.

Related Documentation

- [Tracing L2TP Operations for Subscriber Access on page 371](#)

CHAPTER 36

Contacting Juniper Networks Technical Support

- [Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support on page 375](#)

Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support

Problem **Description:** When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper Networks Technical Support. This topic shows you the most useful logs for a variety of network implementations. In addition to the relevant log information, you must also collect standard troubleshooting information and send it to Juniper Networks Technical Support in your request for assistance.

Solution To collect standard troubleshooting information:

- Redirect the command output to a file.

```
user@host> request support information | save rsi-1
```

To configure logging to assist Juniper Networks Technical Support:

1. Review the following blocks of statements to determine which apply to your configuration.

```
[edit]
set system syslog archive size 100m files 25
set system auto-configuration traceoptions file filename
set system auto-configuration traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions level all
set protocols ppp-service traceoptions flag all
set protocols ppp traceoptions file filename size 100m files 25
set protocols ppp traceoptions level all
set protocols ppp traceoptions flag all
set protocols ppp monitor-session all
set interfaces pp0 traceoptions flag all
set demux traceoptions file filename size 100m files 25
set demux traceoptions level all
set demux traceoptions flag all
set system processes dhcp-service traceoptions file filename
set system processes dhcp-service traceoptions file size 100m
set system processes dhcp-service traceoptions file files 25
set system processes dhcp-service traceoptions flag all
set class-of-service traceoptions file filename
set class-of-service traceoptions file size 100m
set class-of-service traceoptions flag all
set class-of-service traceoptions file files 25
set routing-options traceoptions file filename
set routing-options traceoptions file size 100m
set routing-options traceoptions flag all
set routing-options traceoptions file files 25
set interfaces traceoptions file filename
set interfaces traceoptions file size 100m
set interfaces traceoptions flag all
set interfaces traceoptions file files 25
set system processes general-authentication-service traceoptions file filename
set system processes general-authentication-service traceoptions file size 100m
set system processes general-authentication-service traceoptions flag all
set system processes general-authentication-service traceoptions file files 25
```

2. Copy the relevant statements into a text file and modify the log filenames as you want.
3. Copy the statements from the text file and paste them into the CLI on your router to configure logging.
4. Commit the logging configuration to begin collecting information.



NOTE: The maximum file size for DHCP local server and DHCP relay log files is 1 GB. The maximum number of log files for DHCP local server and DHCP relay is 1000.



BEST PRACTICE: Enable these logs only to collect information when troubleshooting specific problems. Enabling these logs during normal operations can result in reduced system performance.

**Related
Documentation**

- *Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support*

PART 8

Configuration Statements and Operational Commands

- Configuration Statements on page 381
- Operational Commands on page 661

CHAPTER 37

Configuration Statements

- [aaa-access-profile \(L2TP LNS\) on page 387](#)
- [aaa-context \(AAA Options\) on page 388](#)
- [aaa-options \(Access Profile\) on page 389](#)
- [aaa-options \(PPP Profile\) on page 390](#)
- [access \(Dynamic Access Routes\) on page 392](#)
- [access-internal \(Dynamic Access-Internal Routes\) on page 393](#)
- [access-line-information \(L2TP\) on page 394](#)
- [access-profile \(AAA Options\) on page 395](#)
- [address \(L2TP Destination\) on page 396](#)
- [address \(L2TP Tunnel Destination\) on page 397](#)
- [address \(LNS Local Gateway\) on page 398](#)
- [address \(Tunnel Profile Remote Gateway\) on page 398](#)
- [address \(Tunnel Profile Source Gateway\) on page 399](#)
- [address-change-immediate-update on page 399](#)
- [aggregated-inline-services-options \(Aggregated Inline Services\) on page 400](#)
- [allow-snooped-clients on page 401](#)
- [always-write-option-82 on page 402](#)
- [anchor-point \(Pseudowire Subscriber Interfaces\) on page 403](#)
- [assignment-id-format \(L2TP LAC\) on page 404](#)
- [authentication \(Static and Dynamic PPP\) on page 405](#)
- [avp \(L2TP Tunnel Switching\) on page 406](#)
- [bandwidth \(Inline Services\) on page 407](#)
- [bandwidth \(Tunnel Services\) on page 408](#)
- [bearer-type \(L2TP Tunnel Switching\) on page 409](#)
- [bfd on page 410](#)
- [calling-number \(L2TP Tunnel Switching\) on page 411](#)
- [challenge-length \(Static and Dynamic PPP\) on page 412](#)
- [chap on page 413](#)

- chap (Dynamic PPP) on page 414
- chap (L2TP) on page 414
- cisco-nas-port-info (L2TP Tunnel Switching) on page 415
- client on page 416
- delimiter (Access Profile) on page 418
- destination (L2TP) on page 419
- destination-equal-load-balancing (L2TP LAC) on page 420
- destruct-timeout (L2TP) on page 421
- detection-time on page 422
- device-count (Pseudowire Subscriber Interfaces) on page 423
- dhcp-local-server on page 424
- dhcp-relay on page 431
- dhcpv6 (DHCP Local Server) on page 443
- dhcpv6 (DHCP Relay Agent) on page 447
- dial-options on page 453
- dial-options (Dynamic Profiles) on page 454
- disable-calling-number-avp (L2TP LAC) on page 455
- disable-failover-protocol (L2TP) on page 456
- drain on page 457
- dual-stack-group (DHCP Local Server) on page 458
- dual-stack-group (DHCP Relay Agent) on page 460
- duplicate-clients (DHCPv6 Local Server and Relay Agent) on page 462
- duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent) on page 464
- dynamic-profile (L2TP) on page 465
- dynamic-profile (PPP) on page 466
- enable-ipv6-services-for-lac (L2TP) on page 467
- enable-snmp-tunnel-statistics (L2TP) on page 468
- encapsulation (Logical Interface) on page 469
- enforce-strict-scale-limit-license (Subscriber Management) on page 473
- equals (Dynamic Profile) on page 473
- failover-resync on page 474
- failover-within-preference (L2TP LAC) on page 475
- failure-action on page 476
- flexible-vlan-tagging on page 477
- forward-snooped-clients (DHCP Local Server) on page 478
- forward-snooped-clients (DHCP Relay Agent) on page 479
- fpc (MX Series 3D Universal Edge Routers) on page 480

- [gateway-name \(LNS Local Gateway\)](#) on page 481
- [gateway-name \(Tunnel Profile Remote Gateway\)](#) on page 482
- [gateway-name \(Tunnel Profile Source Gateway\)](#) on page 482
- [gres-route-flush-delay \(Subscriber Management\)](#) on page 483
- [group \(DHCP Local Server\)](#) on page 484
- [group \(DHCP Relay Agent\)](#) on page 487
- [group-profile \(Group Profile\)](#) on page 491
- [hierarchical-scheduler \(Subscriber Interfaces on MX Series Routers\)](#) on page 493
- [holddown-interval](#) on page 494
- [hello-interval \(L2TP\)](#) on page 495
- [identification \(Tunnel Profile\)](#) on page 495
- [idle-timeout \(Access\)](#) on page 496
- [idle-timeout \(L2TP\)](#) on page 497
- [ignore-magic-number-mismatch \(Access Group Profile\)](#) on page 498
- [ignore-magic-number-mismatch \(Dynamic Profiles\)](#) on page 500
- [initiate-ncp \(Dynamic and Static PPP\)](#) on page 502
- [inline-services \(FPC Level\)](#) on page 503
- [inline-services \(PIC level\)](#) on page 504
- [input-hierarchical-policer](#) on page 505
- [interface \(Dynamic Routing Instances\)](#) on page 505
- [interface \(L2TP Service Interfaces\)](#) on page 506
- [interface-id](#) on page 507
- [interfaces \(Static and Dynamic Subscribers\)](#) on page 508
- [ip-address-change-notify](#) on page 513
- [ip-reassembly](#) on page 514
- [ip-reassembly \(L2TP\)](#) on page 515
- [ip-reassembly-rules \(Service Set\)](#) on page 516
- [ipcp-suggest-dns-option](#) on page 517
- [keepalive](#) on page 518
- [keepalives](#) on page 519
- [keepalives \(Dynamic Profiles\)](#) on page 520
- [l2tp](#) on page 521
- [l2tp \(Profile\)](#) on page 524
- [l2tp-access-profile](#) on page 525
- [l2tp-maximum-session \(Service Interfaces\)](#) on page 526
- [layer2-liveness-detection \(Receive\)](#) on page 527
- [layer2-liveness-detection \(Send\)](#) on page 528

- [lcp-renegotiation](#) on page 530
- [liveness-detection](#) on page 531
- [local-gateway \(L2TP LNS\)](#) on page 532
- [lockout-timeout \(L2TP Destination Lockout\)](#) on page 533
- [logical-system \(Tunnel Profile\)](#) on page 534
- [mac](#) on page 534
- [mac-address \(Dynamic Access-Internal Routes\)](#) on page 535
- [match-direction \(IP Reassembly Rule\)](#) on page 536
- [maximum-sessions \(L2TP\)](#) on page 537
- [maximum-sessions-per-tunnel](#) on page 538
- [max-sessions \(Tunnel Profile\)](#) on page 539
- [medium \(Tunnel Profile\)](#) on page 539
- [method](#) on page 540
- [metric \(Dynamic Access-Internal Routes\)](#) on page 542
- [minimum-interval](#) on page 543
- [minimum-receive-interval](#) on page 544
- [minimum-retransmission-timeout \(L2TP Tunnel\)](#) on page 545
- [mtu](#) on page 546
- [multiplier](#) on page 550
- [name \(L2TP Destination\)](#) on page 551
- [name \(L2TP Tunnel Destination\)](#) on page 552
- [no-adaptation](#) on page 553
- [nas-port-method \(L2TP LAC\)](#) on page 554
- [nas-port-method \(Tunnel Profile\)](#) on page 554
- [next-hop \(Dynamic Access-Internal Routes\)](#) on page 555
- [next-hop-service](#) on page 556
- [no-allow-snooped-clients](#) on page 557
- [no-gratuitous-arp-request](#) on page 558
- [no-snoop \(DHCP Local Server and Relay Agent\)](#) on page 559
- [no-vlan-id-validate](#) on page 560
- [on-demand-ip-address](#) on page 561
- [overrides \(DHCP Relay Agent\)](#) on page 562
- [overrides \(Enhanced Subscriber Management\)](#) on page 564
- [override-result-code \(L2TP Profile\)](#) on page 565
- [pap](#) on page 566
- [pap \(Dynamic PPP\)](#) on page 567
- [pap \(L2TP\)](#) on page 567

- [parse-direction \(Access Profile\) on page 568](#)
- [pic \(M Series and T Series Routers\) on page 569](#)
- [pool \(L2TP Service Interfaces\) on page 570](#)
- [pp0 \(Dynamic PPPoE\) on page 571](#)
- [ppp \(Group Profile\) on page 573](#)
- [ppp-options on page 574](#)
- [ppp-options \(Dynamic PPP\) on page 576](#)
- [ppp-options \(L2TP\) on page 578](#)
- [preference \(Subscriber Management\) on page 579](#)
- [preference \(Tunnel Profile\) on page 580](#)
- [primary-interface \(Aggregated Inline Services\) on page 581](#)
- [profile \(Access\) on page 582](#)
- [proxy-mode on page 587](#)
- [ps0 \(Pseudowire Subscriber Interfaces\) on page 588](#)
- [pseudowire-service \(Pseudowire Subscriber Interfaces\) on page 589](#)
- [qualified-next-hop \(Subscriber Management\) on page 590](#)
- [reject-unauthorized-ipv6cp on page 591](#)
- [relay-option-82 on page 592](#)
- [remote-gateway \(Tunnel Profile\) on page 593](#)
- [report-ingress-shaping-rate \(Dynamic CoS Interfaces\) on page 594](#)
- [request services l2tp destination unlock](#)
- [retransmission-count-established \(L2TP\) on page 596](#)
- [retransmission-count-not-established \(L2TP\) on page 597](#)
- [route \(Access\) on page 598](#)
- [route \(Access Internal\) on page 599](#)
- [route-suppression \(DHCP Local Server and Relay Agent\) on page 600](#)
- [routing-instance \(Tunnel Profile\) on page 601](#)
- [routing-instance \(L2TP Destination\) on page 601](#)
- [routing-instance \(L2TP Tunnel Destination\) on page 602](#)
- [routing-instances \(Dynamic Profiles\) on page 603](#)
- [routing-options \(Dynamic Profiles\) on page 605](#)
- [rule \(IP Reassembly\) on page 606](#)
- [rx-connect-speed-when-equal \(L2TP LAC\) on page 607](#)
- [rx-window-size \(L2TP\) on page 608](#)
- [secondary-interface \(Aggregated Inline Services\) on page 609](#)
- [secret \(Tunnel Profile\) on page 610](#)
- [service-device-pool \(L2TP\) on page 610](#)


- [service-device-pools \(L2TP Service Interfaces\)](#) on page 611
- [service-interface \(L2TP Processing\)](#) on page 612
- [service-profile \(L2TP\)](#) on page 613
- [service-rate-limiter \(Access\)](#) on page 615
- [session-mode](#) on page 616
- [session-options](#) on page 617
- [sessions-limit-group \(L2TP\)](#) on page 618
- [sessions-limit-group \(L2TP Client Profile\)](#) on page 619
- [shared-secret](#) on page 619
- [soft-gre](#) on page 620
- [source-gateway \(Tunnel Profile\)](#) on page 621
- [stacked-vlan-tagging](#) on page 622
- [statistics \(Access Profile\)](#) on page 622
- [strip-user-name \(Access Profile\)](#) on page 623
- [subscriber-context \(AAA Options\)](#) on page 624
- [subscriber-management \(Subscriber Management\)](#) on page 625
- [tag \(Access\)](#) on page 626
- [threshold \(detection-time\)](#) on page 627
- [threshold \(transmit-interval\)](#) on page 628
- [tos-reflect \(L2TP\)](#) on page 629
- [trace \(DHCP Relay Agent\)](#) on page 630
- [traceoptions \(Services L2TP\)](#) on page 631
- [traceoptions \(Protocols PPP Service\)](#) on page 635
- [traceoptions \(Subscriber Management\)](#) on page 638
- [transmit-interval](#) on page 639
- [tunnel \(L2TP\)](#) on page 640
- [tunnel \(Tunnel Profile\)](#) on page 641
- [tunnel-group](#) on page 642
- [tunnel-profile \(L2TP Tunnel Switching\)](#) on page 643
- [tunnel-profile \(Tunnel Profile\)](#) on page 644
- [tunnel-switch-profile \(L2TP Tunnel Switching, Application\)](#) on page 645
- [tunnel-switch-profile \(L2TP Tunnel Switching, Definition\)](#) on page 646
- [tx-address-change \(L2TP LAC\)](#) on page 647
- [tx-connect-speed-method \(L2TP LAC\)](#) on page 648
- [type \(Tunnel Profile\)](#) on page 651
- [unit \(Dynamic PPPoE\)](#) on page 652
- [untagged](#) on page 653

- [user-group-profile](#) on page 654
- [version \(BFD\)](#) on page 655
- [weighted-load-balancing \(L2TP LAC\)](#) on page 656
- [vlan-id \(Dynamic Profiles\)](#) on page 657
- [vlan-tagging](#) on page 658
- [vlan-tagging \(Dynamic\)](#) on page 659
- [vlan-tags](#) on page 660

aaa-access-profile (L2TP LNS)

Syntax	<code>aaa-access-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group name],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]</code>
Release Information	Statement introduced in Junos OS Release 11.4. Support at the <code>[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]</code> hierarchy level introduced in Junos OS Release 12.1.
Description	Specify a AAA access profile that overrides the AAA access profile configured for the routing instance with the access-profile statement. You can configure a profile to specify the RADIUS server settings for a tunnel group or for a LAC client, or both. The AAA access profile configured for the client takes precedence over the AAA access profile configured for the tunnel group, which takes precedence over the access profile configured for the routing instance.
Options	<i>profile-name</i> —Name of the local access profile for the tunnel group or client.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring an L2TP LNS with Inline Service Interfaces on page 239 • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274 • Configuring an L2TP Access Profile on the LNS on page 246


aaa-context (AAA Options)

Syntax	<code>aaa-context <i>aaa-context-name</i>;</code>
Hierarchy Level	[edit access aaa-options <i>aaa-options-name</i>]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	Specify the logical-system:routing-instance (LS:RI) that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting. For example, this may correspond to the LS:RI for a retail ISP that provides services to the subscriber.
<div> NOTE: Only the default (master) logical system is supported.</div>	
Options	<i>aaa-context-name</i> —Name of the logical-system:routing-instance.
Required Privilege Level	access—To view this statement in the configuration. access-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Session Options for Subscriber Access• Configuring Username Modification for Subscriber Sessions on page 179• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

aaa-options (Access Profile)

Syntax	<pre>aaa-options <i>aaa-options-name</i> { <i>aaa-context</i> <i>aaa-context-name</i>; <i>access-profile</i> <i>profile-name</i>; <i>subscriber-context</i> <i>subscriber-context-name</i> }</pre>
Hierarchy Level	[edit access]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	Define a set of AAA options for authorizing and configuring a subscriber or set of subscribers with a subscriber access profile.
Options	<p><i>aaa-options-name</i>—Name of the set of options.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Session Options for Subscriber Access • Configuring Username Modification for Subscriber Sessions on page 179 • Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245 • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

aaa-options (PPP Profile)

Syntax	<code>aaa-options <i>aaa-options-name</i>;</code>
Hierarchy Level	<code>[edit access group-profile <i>profile-name</i> ppp aaa-options],</code> <code>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" aaa-options]</code>
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify a set of AAA options that is used for authentication of PPP subscribers. The set of options is defined globally with the <code>aaa-options <i>aaa-options-name</i></code> statement at the <code>[edit access]</code> hierarchy level.</p> <p>You can specify the option set in a dynamic PPP profile or in a group profile.</p> <ul style="list-style-type: none"> In a dynamic PPP profile—In this case, usernames are examined and modified for dynamic PPP subscribers logging in by means of the subscriber and AAA contexts that are specified in the AAA options set. The option set must include the <code>access-profile <i>profile-name</i></code> statement to specify the name of a subscriber access profile. In a group profile—In this case, usernames are examined and modified for tunneled PPP subscribers on the LNS logging in by means of the subscriber and AAA contexts that are specified in the AAA options set.
	<p> NOTE: When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.</p> <p>This meant that <code>aaa-options</code> configured in a group profile is not applied when the dynamic profile includes any PPP-option, even when the dynamic profile does not include <code>aaa-options</code>.</p>
Options	<code><i>aaa-options-name</i></code> —Name of the set of options.
Required Privilege Level	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding Session Options for Subscriber Access Configuring Username Modification for Subscriber Sessions on page 179 Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245

- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)

access (Dynamic Access Routes)

Syntax

```
access {
  route prefix {
    next-hop next-hop;
    metric route-cost;
    preference route-distance;
    tag route-tag;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance **routing-options**],
[edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance **routing-options** rib *routing-table-name*],
[edit dynamic-profiles *profile-name* **routing-options**]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options] and [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options rib *routing-table-name*] hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure access routes.



NOTE: Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If you configure the access-internal statement in the dynamic profile, it is ignored. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute (22) or Framed-IPv6-Route attribute [99].

Options The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

access-internal (Dynamic Access-Internal Routes)

Syntax

```
access-internal {
  route subscriber-ip-address {
    qualified-next-hop underlying-interface {
      mac-address address;
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance *routing-options*],
[edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance *routing-options* rib *routing-table-name*],
[edit dynamic-profiles *routing-options*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance *routing-options*] and [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance *routing-options* rib *routing-table-name*] hierarchy levels introduced in Junos OS Release 10.1.

Description (Releases earlier than Junos OS Release 15.1) Dynamically configure access-internal routes. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.



NOTE: Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If you configure the access-internal statement in the dynamic profile, it is ignored. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute (22) or Framed-IPv6-Route attribute [99].

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

- Related Documentation**
- [Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32](#)
 - [Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 123](#)

access-line-information (L2TP)

Syntax	<code>access-line-information <connection-speed-update>;</code>
Hierarchy Level	<code>[edit services l2tp]</code> , <code>[edit services l2tp destination ip-address]</code>
Release Information	Statement introduced in Junos OS Release 14.1. Support at the <code>[edit services l2tp]</code> hierarchy level introduced in Junos OS Release 14.2. Support for the LNS introduced in Junos OS Release 17.4R1 on MX Series routers.
Description	<p>Configure a LAC to forward subscriber line identification and other DSL attributes in ICRQ messages to the LNS by means of L2TP AVPs for all tunnels to all LNSs or for only tunnels with the specified endpoint for a particular LNS. Optionally, configure the LAC to send initial line rates in ICCN messages and subsequent rate updates in CSUN messages.</p> <p>Configure an LNS to process such line information for all tunnels from all LACs or for only tunnels with the specified endpoint for a particular LAC. Optionally, configure the LNS to process rate updates received in CSUN messages from the LAC.</p> <p>Including this statement at the <code>[edit services l2tp destination ip-address]</code> hierarchy level is useful when you know that some endpoints in the network do not support this feature or have an incorrect implementation. Configuring at this level enables you to restrict the transmission or processing of this information to only LACs and LNSs, respectively, that you know support the feature.</p> <p>This statement has no effect on existing subscribers; it applies only to new subscribers.</p>
Options	<p>connection-speed-update—(Optional) On the LAC, include the Connect Speed Update Enable AVP (98) in ICCN messages from the LAC to alert the LNS that the LAC might send CSUN messages that report speed changes originating with the ANCP agent.</p> <p>On the LNS, enable processing of CSUN updates. If this option is not configured on the LNS, CSUN updates cannot be processed even when the Connect Speed Update Enable AVP (98) is received from the LAC. In that case, only rates received in AVP 24 (Tx speed) and AVP 38 (Rx speed) in ICCN messages can be applied.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	• Configuring the Reporting and Processing of Subscriber Access Line Information on page 230

access-profile (AAA Options)

Syntax	<code>access-profile <i>profile-name</i>;</code>
Hierarchy Level	[edit access aaa-options <i>aaa-options-name</i>]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	Specify the name of the access profile that includes the username stripping configuration.
Options	<i>profile-name</i> —Name of the subscriber access profile that includes a subscriber username stripping configuration.
Required Privilege Level	<code>access</code> —To view this statement in the configuration. <code>access-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Session Options for Subscriber Access</i>• <i>Configuring Subscriber Session Timeout Options</i>• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

address (L2TP Destination)

Syntax address *ip-address* {
 access-line-information <connection-speed-update>;
 drain;
 routing-instance *routing-instance-name* {
 drain;
 }
 }

Hierarchy Level [edit services l2tp destination]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the IP address and other attributes for the L2TP destination.

Options *ip-address*—IP address of the destination; corresponds to the IP address that is used by LACs to identify the LNS.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring L2TP Drain on page 166](#)

address (L2TP Tunnel Destination)

Syntax	<pre> address <i>ip-address</i>; { drain; routing-instance <i>routing-instance-name</i> { drain; } } </pre>
Hierarchy Level	[edit services l2tp tunnel name <i>name</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the IP address for the L2TP tunnel destination when the name statement at the [edit services l2tp tunnel] hierarchy level specifies only the name of the tunnel rather than both the name and the destination address. Do not include the address statement when the name statement provides both the tunnel name and the destination address.
Options	<p>ip-address—IP address of the tunnel destination.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring L2TP Drain on page 166

address (LNS Local Gateway)

Syntax	<code>address <i>address</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i> local-gateway]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the local (LNS) IP address for L2TP tunnel.
Options	<i>address</i> —Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Local Gateway Address and PIC.</i>• <i>Configuring L2TP Tunnel Groups</i>• <i>Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274</i>

address (Tunnel Profile Remote Gateway)

Syntax	<code>address <i>server-ip-address</i>;</code>
Hierarchy Level	<code>[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i> remote-gateway]</code>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the IP address of the remote gateway device at the L2TP tunnel endpoint, the LNS.
Options	<i>server-ip-address</i> —IP address of the remote gateway device. Default: 0.0.0.0.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Tunnel Profile for Subscriber Access on page 207</i>

address (Tunnel Profile Source Gateway)

Syntax	<code>address <i>client-ip-address</i>;</code>
Hierarchy Level	<code>[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i> source-gateway]</code>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the IP address of the source gateway device at the local L2TP tunnel endpoint, the LAC. This value overrides the default address for the logical system or routing instance.
Options	<p><i>client-ip-address</i>—IP address of the source gateway device.</p> <p>Default: 0.0.0.0.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Tunnel Profile for Subscriber Access on page 207

address-change-immediate-update

Syntax	<code>address-change-immediate-update;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> accounting]</code>
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>Configure the router to send an Interim-Accounting message to the RADIUS server immediately after on-demand IPv4 allocation and de-allocation.</p> <p>Changes to this setting take effect for new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.</p>
Default	This functionality is disabled by default.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes • Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

aggregated-inline-services-options (Aggregated Inline Services)

Syntax `aggregated-inline-services-options {
 primary-interface interface-name;
 secondary-interface interface-name;
 }`

Hierarchy Level [edit interfaces asix]

Release Information Statement introduced in Junos OS Release 16.2.

Description Configure the members of an aggregated inline service interface bundle to provide 1:1 stateful LNS redundancy for an LNS sessions in a tunnel group.



BEST PRACTICE: Follow these guidelines:

- You must configure unit 0 family inet for each bundle; otherwise, the session fails to come up.
- The primary (active) and secondary (backup) interfaces must be on different MPCs. If you configure both interfaces on the same MPC, the subsequent configuration commit fails.
- The bandwidth configured at the [edit chassis fpc slot pic *number* inline-services bandwidth] hierarchy level must be the same for both member links.
- An si interface configured as a member of an aggregated inline service bundle cannot be configured as a member of another bundle group.
- An si interface configured as a member of an aggregated inline service bundle cannot also be used for any function that is not related to aggregated services; for example, it cannot be used for inline IP reassembly.
- When you configure an si interface as a member of an aggregated inline services bundle, you can no longer configure that si interface independently. You can configure only the parent bundle; the bundle's configuration is applied immediately to all member interfaces.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

allow-snooped-clients

Syntax allow-snooped-clients;

Hierarchy Level [edit forwarding-options dhcp-relay **dhcpv6** group *group-name* interface *interface-name* **overrides**],
 [edit forwarding-options dhcp-relay **dhcpv6** group *group-name* **overrides**],
 [edit forwarding-options dhcp-relay **dhcpv6** **overrides**],
 [edit forwarding-options dhcp-relay group *group-name* interface *interface-name* **overrides**],
 [edit forwarding-options dhcp-relay group *group-name* **overrides**],
 [edit forwarding-options dhcp-relay **overrides**],
 [edit logical-systems *logical-system-name* forwarding-options **dhcp-relay** ...],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options **dhcp-relay** ...],
 [edit routing-instances *routing-instance-name* forwarding-options **dhcp-relay** ...]

Release Information Statement introduced in Junos OS Release 10.2.
 Support at the [edit ... **dhcpv6**] hierarchy levels introduced in Junos OS Release 12.1.

Description Explicitly enable DHCP snooping support on the DHCP relay agent.

Use the statement at the [edit ... **dhcpv6**] hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.

Default DHCP snooping is disabled by default.



NOTE: On EX4300 and EX9200 switches, the **allow-snooped-clients** statement is enabled by default at the [edit forwarding-options dhcp-relay **overrides**] hierarchy level.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [Extended DHCP Relay Agent Overview](#)
- [Overriding the Default DHCP Relay Configuration Settings](#)
- [DHCP Snooping Support on page 41](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48](#)

always-write-option-82

Syntax	<code>always-write-option-82;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay overrides],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i></code> <code> overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code> forwarding-options dhcp-relay overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code> forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group</code> <code> <i>group-name</i> overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group</code> <code> <i>group-name</i> interface <i>interface-name</i> overrides]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:</p> <ul style="list-style-type: none">• If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.• If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Extended DHCP Relay Agent Overview</i>

anchor-point (Pseudowire Subscriber Interfaces)

Syntax	<code>anchor-point <i>lt-device</i>;</code>
Hierarchy Level	<code>[edit interfaces ps0]</code>
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>Specify the anchor-point, a logical tunnel (lt) interface that identifies the logical tunnel interface that terminates the MPLS pseudowire tunnel at the access node. The other end of the tunnel terminates on the pseudowire subscriber logical interface, which is configured on an MX Series router or PTX Series router that hosts subscriber management and enables you to perform subscriber management services at the interface.</p> <p>The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD), and supports a circuit-oriented Layer 2 protocol (either Layer 2 VPN or Layer 2 circuit). The Layer 2 protocol provides the transport and service logical interfaces, and supports the protocol family (IPv4, IPv6, or PPPoE).</p>



NOTE: You cannot dynamically change an anchor point that has active pseudowire devices stacked above it. If you need to change such an anchor point, you must perform the following steps:

1. Deactivate the stacked pseudowires and commit. This may require bringing down any subscribers using the pseudowires.

```
[edit interfaces]
user@host# deactivate psnumber
user@host# commit
```

2. Change the anchor on the deactivated pseudowire and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point lt-new-lt-interface-number
user@host# commit
```


3. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

Options	<i>lt-device</i> —An lt device in the format <i>lt-fpc/pic/port</i>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

- Related Documentation**
- [Pseudowire Subscriber Logical Interfaces Overview on page 307](#)
 - [Configuring a Pseudowire Subscriber Logical Interface on page 313](#)
 - [Configuring a Pseudowire Subscriber Logical Interface Device on page 315](#)

assignment-id-format (L2TP LAC)

Syntax	assignment-id-format (assignment-id client-server-id);
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Set the format for the name used for a tunnel, the tunnel assignment ID.
	<div> NOTE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing <code>no services l2tp tunnel assignment-id-format</code>.</div>
Default	assignment-id
Options	<p>assignment-id—The tunnel name corresponds to RADIUS attribute Tunnel-Assignment-Id [82].</p> <p>client-server-id—The tunnel name is a combination of RADIUS attributes Tunnel-Client-Auth-Id [90], Tunnel-Server-Auth-Id [91], and Tunnel-Assignment-Id [82].</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Setting the Format for the Tunnel Name on page 206

authentication (Static and Dynamic PPP)

Syntax	<code>authentication [<i>authentication-protocols</i>];</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options], [edit interfaces pp0 unit <i>unit-number</i> ppp-options]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	<p>Specify the order in which the router tries to negotiate PPP authentication protocols when verifying that a PPP client can access the network. By default, the router tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication first, and then tries Password Authentication Protocol (PAP) authentication if the attempt to negotiate CHAP authentication is unsuccessful.</p> <p>You can specify one or both authentication protocols. If you specify both CHAP and PAP in either order, you must enclose the set of protocol names within square brackets ([]).</p>
Options	<p><i>authentication-protocols</i>—One or both of the following PPP authentication protocols:</p> <ul style="list-style-type: none"> • chap—Challenge Handshake Authentication Protocol • pap—Password Authentication Protocol
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Controlling the Negotiation Order of PPP Authentication Protocols on page 132

avp (L2TP Tunnel Switching)

Syntax	<pre>avp { bearer-type; calling-number; cisco-nas-port-info; }</pre>
Hierarchy Level	[edit access tunnel-switch-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	<p>Specify the action taken on L2TP AVPs that are negotiated when the first session is created; these AVPs are contained in the L2TP packets that are switched by the tunnel switch profile.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Tunnel Switching on page 161

bandwidth (Inline Services)

Syntax	<code>bandwidth (1g 10g);</code>
Hierarchy Level	[edit chassis <code>fpc slot-number pic number inline-services</code>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.
Options	<p>1g—Reserves 1 Gbps of bandwidth for tunnel traffic. Configuring a bandwidth of 1 Gbps creates a virtual tunnel interface that is represented as <code>si-<fpc/pic/port></code>.</p> <p>10g—Reserves 10 Gbps of bandwidth for tunnel traffic. Configuring a bandwidth of 10 Gbps creates a virtual tunnel interface that is represented as <code>si-<fpc/pic/port></code>.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Inline Service Interfaces on page 251 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

bandwidth (Tunnel Services)

Syntax	<code>bandwidth <i>bandwidth-value</i>;</code>
Hierarchy Level	<code>[edit chassis fpc <i>slot-number</i> pic <i>number</i> tunnel-services]</code>
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3X54 for ACX Series routers.
Description	(ACX Series, MX Series 3D Universal Edge Routers and T4000 Core Routers only) Specify the amount of bandwidth in gigabits per second to reserve for tunnel services. For ACX Series routers, you can configure a bandwidth of only 1 Gbps and 10 Gbps for logical tunnel (lt-) interfaces.
Options	<i>bandwidth-value</i> —Amount of bandwidth in Gbps to reserve for tunnel services. On MX Series routers, the bandwidth values can be 1g and multiples of 10g up to 100g . On T4000 routers, the bandwidth values are multiples of 10g up to 100g .



NOTE: The bandwidth that you specify determines the port number of the tunnel interfaces that are created. When you specify a bandwidth of **1g**, the port number is always 10. When you specify any other bandwidth, the port number is always 0.



NOTE: If you specify a bandwidth that is not compatible with the type of DPCs or MPCs and their respective Packet Forwarding Engine, tunnel services are not activated. For example, you cannot specify 1 gigabit per second bandwidth for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC or 16x10GE 3D MPC.

When the tunnel bandwidth is unspecified in the Routing Engine CLI, the maximum tunnel bandwidth for MPC3E is 60G.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC</i> • <i>Tunnel Interface Configuration on MX Series Routers Overview</i> • <i>Configuring Tunnel Interfaces on T4000 Routers</i> • <i>Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC</i> • <i>Example: Configuring Tunnel Interfaces on the MPC3E</i>
------------------------------	--

- *tunnel-services (Chassis)*

bearer-type (L2TP Tunnel Switching)

Syntax	<code>bearer-type <i>action</i>;</code>
Hierarchy Level	[edit access tunnel-switch-profile <i>profile-name</i> avp]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the action taken on the Bearer Type AVP (18) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.
Options	<p><i>action</i>—One of the following actions:</p> <ul style="list-style-type: none"> • drop—Drop the AVP. • regenerate—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session. • relay—Forward the AVP transparently as is and send it in the switched packet. <p>Default: relay</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring L2TP Tunnel Switching on page 161


bfd

Syntax	<pre> bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } </pre>
Hierarchy Level	<pre> [edit system services dhcp-local-server liveness-detection <i>method</i>], [edit system services dhcp-local-server dhcpv6 liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection <i>method</i>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection <i>method</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection <i>method</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89


calling-number (L2TP Tunnel Switching)

Syntax	calling-number <i>action</i> ;
Hierarchy Level	[edit access tunnel-switch-profile <i>profile-name</i> avp]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the action taken on the Calling Number AVP (22) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.
Options	<p>action—One of the following actions:</p> <ul style="list-style-type: none">• drop—Drop the AVP.• regenerate—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.• relay—Forward the AVP transparently as is and send it in the switched packet. <p>Default: relay</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Tunnel Switching on page 161

challenge-length (Static and Dynamic PPP)

Syntax	challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i> ;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options chap], [edit interfaces pp0 unit <i>unit-number</i> ppp-options chap]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Modify the length of the Challenge Handshake Authentication Protocol (CHAP) challenge by specifying the minimum and maximum allowable length, in bytes.
<div> BEST PRACTICE: We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.</div>	
Options	<p><i>minimum-length</i>—Minimum length, in bytes, of the CHAP challenge.</p> <p>Range: 8 through 63</p> <p>Default: 16</p> <p><i>maximum-length</i>—Maximum length, in bytes, of the CHAP challenge. The <i>maximum-length</i> must be equal to or greater than the <i>minimum-length</i>.</p> <p>Range: 8 through 63</p> <p>Default: 32</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">• Modifying the CHAP Challenge Length on page 127

chap

Syntax	<pre>chap { access-profile <i>name</i>; challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>; default-chap-secret <i>name</i>; local-name <i>name</i>; passive; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> ppp-options],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Allow each side of a link to challenge its peer, using a “secret” known only to the authenticator and that peer. The secret is not sent over the link.</p> <p>By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.</p> <p>For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> • atm-ppp-llc—PPP over AAL5 LLC encapsulation. • atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation. <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;">  <p>BEST PRACTICE: On inline service (si) interfaces for L2TP, only the chap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.</p> </div> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the PPP Challenge Handshake Authentication Protocol • Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245 • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

chap (Dynamic PPP)

Syntax	<pre>chap { challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>; local-name <i>name</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options] hierarchy level introduced in Junos OS Release 12.2.
Description	Specify CHAP authentication in a PPP dynamic 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">• Dynamic Profiles Overview• Configuring Dynamic Authentication for PPP Subscribers on page 125• Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

chap (L2TP)

Syntax	<pre>chap;</pre>
Hierarchy Level	[edit access group-profile <i>profile-name</i> ppp ppp-options]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	(MX Series routers only) Specify CHAP authentication for PPP subscribers in an L2TP LNS user group profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Configuring an L2TP LNS with Inline Service Interfaces on page 239

cisco-nas-port-info (L2TP Tunnel Switching)

Syntax	<code>cisco-nas-port-info action;</code>
Hierarchy Level	[edit access tunnel-switch-profile <i>profile-name</i> avp]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	<p>Define a tunnel profile for subscriber access.</p> <p>Specify the action taken on the Cisco NAS Port Info AVP (100) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.</p>
Options	<p>action—One of the following actions:</p> <ul style="list-style-type: none">• drop—Drop the AVP.• regenerate—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.• relay—Forward the AVP transparently as is and send it in the switched packet. <p>Default: relay</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Tunnel Switching on page 161

client

```

Syntax  client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
        allowed-proxy-pair {
            remote remote-proxy-address local local-proxy-address;
        }
        pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
        ike-policy policy-name;
        interface-id string-value;
    }
    l2tp {
        aaa-access-profile profile-name;
        interface-id interface-id;
        lcp-renegotiation;
        local-chap;
        maximum-sessions number;
        maximum-sessions-per-tunnel number;
        multilink {
            drop-timeout milliseconds;
            fragment-threshold bytes;
        }
        override-result-code session-out-of-resource;
        ppp-authentication (chap | pap);
        ppp-profile profile-name;
        sessions-limit-group;
        service-profile profile-name (parameter) & profile-name;
        shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
        cell-overhead;
        encapsulation-overhead bytes;
        framed-ip-address ip-address;
        framed-pool framed-pool;
        idle-timeout seconds;
        interface-id interface-id;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
    }
    user-group-profile profile-name;
}

```

Hierarchy Level [edit access [profile](#) *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the peer identity.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *client-name*—A peer identity. For L2TP clients, you can use a special name to configure a default client. This client enables the LNS to accept any LAC to establish the session. On M Series routers, use * for the default client configuration. On MX Series routers, use **default**.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the L2TP Client*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- [Configuring an L2TP Access Profile on the LNS on page 246](#)

delimiter (Access Profile)

Syntax	<code>delimiter <i>delimiter</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> session-options strip-user-name]</code>
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify up to eight characters that the router uses to determine which part of the subscriber login string to discard—leaving the remainder for use as a modified username—when subscriber username stripping is configured in a subscriber access profile. The characters to the right of the delimiter are discarded along with the delimiter. Use the parse-direction statement when more than one delimiter appears in a username to determine the characters that are stripped by identifying the desired delimiter. A given subscriber login string can result in multiple different modified usernames depending on the number and placement of delimiters and the direction of stripping.</p>
Default	None. You must always configure a delimiter.
Options	<i>delimiter</i> —Character that specifies the boundary between the part of the original username that is kept and the part that is discarded.
Required Privilege Level	<code>access</code> —To view this statement in the configuration. <code>access-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Session Options for Subscriber Access• Configuring Username Modification for Subscriber Sessions on page 179• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

destination (L2TP)

Syntax

```
destination {
  address ip-address {
    access-line-information <connection-speed-update>;
    drain;
    routing-instance routing-instance-name {
      drain;
    }
  }
  lockout-timeout seconds;
  name destination-name {
    drain;
  }
}
```

Hierarchy Level [edit services [l2tp](#)]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure attributes for all L2TP destinations or a specified L2TP destination.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- [Configuring the L2TP Destination Lockout Timeout on page 165](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring the Reporting and Processing of Subscriber Access Line Information on page 230](#)

destination-equal-load-balancing (L2TP LAC)

Syntax	destination-equal-load-balancing;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Enable the LAC to balance the L2TP session load equally across multiple LNSs by selecting tunnels according to how many sessions currently exist for the destination and tunnel.</p> <p>Disabled by default. By default, tunnel selection within a preference level is strictly random. The weighted-load-balancing statement must be disabled to successfully enable this statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions on page 211• Configuring the L2TP LAC Tunnel Selection Parameters on page 209• LAC Tunnel Selection Overview on page 183

destruct-timeout (L2TP)

Syntax	<code>destruct-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Set how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
<div>  <p>BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing <code>no services l2tp destruct-timeout</code>.</p> </div>	
Options	<p><i>seconds</i>—Length of the destruct timeout.</p> <p>Range: 10 through 3600</p> <p>Default: 300</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Setting the L2TP Destruct Timeout on page 164 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

detection-time

Syntax	<pre>detection-time { <i>threshold</i> <i>milliseconds</i>; }</pre>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

device-count (Pseudowire Subscriber Interfaces)

Syntax	device-count <i>number</i> ;
Hierarchy Level	[edit chassis pseudowire-service]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Configure the number of pseudowire logical devices available to the router.
Options	<i>number</i> —Number of devices. Range: 1 through 2048
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Pseudowire Subscriber Logical Interfaces Overview on page 307• Configuring a Pseudowire Subscriber Logical Interface on page 313• Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 315

dhcp-local-server

```
Syntax  dhcp-local-server {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            access-profile profile-name;
            authentication {
                ...
            }
            duplicate-clients incoming-interface;
            group group-name {
                access-profile profile-name;
                authentication {
                    ...
                }
            }
            interface interface-name {
                access-profile profile-name;
                exclude;
                overrides {
                    asymmetric-lease-time seconds;
                    asymmetric-prefix-lease-time seconds;
                    interface-client-limit number;
                    multi-address-embedded-option-response;
                    process-inform {
                        pool pool-name;
                    }
                    protocol-attributes attribute-set-name;
                    rapid-commit;
                }
                service-profile dynamic-profile-name;
                trace;
                upto upto-interface-name;
            }
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                    }
                }
            }
        }
    }
```

```

    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
}
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {

```

```

        max-consecutive-retries number;
        transmit-interval interval;
    }
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    delegated-pool;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    support-option-pd-exclude;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
reauthenticate (<lease-renewal> <remote-id-mismatch >);
requested-ip-network-match subnet-mask;
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
}
dual-stack-group name {
    access-profile access-profile;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
}

```

```

classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
}
dual-stack-interface-client-limit number;
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
on-demand-address-allocation;
protocol-master (inet | inet6);
reauthenticate (<lease-renewal> <remote-id-mismatch>);
service-profile service-profile;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    authentication {
        ...
    }
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
        primary-profile-name>;
    interface interface-name {
        exclude;
        overrides {
            asymmetric-lease-time seconds;
            client-discover-match (option60-and-option82 | incoming-interface);
            include-option-82 {
                forcerenew;
                nak;
            }
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    }
}

```

```

method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
  layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
  }
}
}
overrides {
  asymmetric-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  include-option-82 {
    forcerenew;
    nak;
  }
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
}
requested-ip-network-match subnet-mask
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
    }
  }
}

```



```

    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
  layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
  }
}
on-demand-address-allocation;
overrides {
  asymmetric-lease-time seconds;
  client-discover-match <option60-and-option82 | incoming-interface>;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
protocol-master;
reauthenticate (<lease-renewal> <remote-id-mismatch >);
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],
 [edit logical-systems *logical-system-name* system services],
 [edit routing-instances *routing-instance-name* system services],
 [edit system services]

Release Information Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch to enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and, therefore, is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcpv6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.



NOTE: When you configure the **dhcp-local-server** statement at the routing instance hierarchy level, you must use a routing instance type of **virtual-router**.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Extended DHCP Local Server Overview*
- *DHCPv6 Local Server Overview*

dhcp-relay

```
Syntax  dhcp-relay {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
        trigger automatic;
    }
    dhcpv6 {
        access-profile profile-name;
        active-server-group server-group-name;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
        trigger automatic;
    }
    duplicate-clients incoming-interface;
```

```
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
  non-configured-interfaces);
group group-name {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name interface-name;
      logical-system-name;
      mac-address mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
}
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
interface interface-name {
  access-profile profile-name;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  exclude;
  overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
```

```

    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
}
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
relay-agent-interface-id {
    include-irb-and-l2;
}

```

```
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
}
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
    }
}
```

```

    }
    detection-time {
        threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
route-suppression;
service-profile dynamic-profile-name;
}
no-snoop;
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
}

```

```

        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
relay-option-vendor-specific{
    host-name;
    location;
    remote-id-mismatch disconnect;
    route-suppression;
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
}
dual-stack-group dual-stack-group-name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
}

```



```

    }
}
protocol-master;
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 [circuit-id] [remote-id];
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
}
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}

```

```
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
interface interface-name {
  access-profile profile-name;
  exclude;
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
}
overrides {
  allow-no-end-option;
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  client-discover-match <option60-and-option82 | incoming-interface>;
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
```

```

}
overrides {
  allow-no-end-option
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  server-id-override
}
remote-id-mismatch disconnect;
route-suppression:
service-profile dynamic-profile-name;

```

```
}
leasequery {
  attempts number-of-attempts;
  timeout seconds;
}
lease-time-validation {
  lease-time-threshold seconds;
  violation-action action;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
  }
}
no-snoop;
overrides {
  allow-no-end-option
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
```

```

}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  server-id-override
}
}
remote-id-mismatch disconnect;
route-suppression:
server-group {
  server-group-name {
    server-ip-address;
  }
}
server-response-time seconds;
service-profile dynamic-profile-name;

```

Hierarchy Level [edit forwarding-options],
 [edit logical-systems *logical-system-name* forwarding-options],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options],
 [edit routing-instances *routing-instance-name* forwarding-options]

Release Information Statement introduced in Junos OS Release 8.3.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*
- *DHCP Relay Proxy Overview*
- *Using External AAA Authentication Services with DHCP*

dhcpcv6 (DHCP Local Server)

```
Syntax  dhcpcv6 {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    duplicate-clients incoming-interface;
    group group-name {
        access-profile profile-name;
        authentication {
            ...
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
    }
}
```

```

    delete-binding-on-renegotiation;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delegated-pool;
    delete-binding-on-renegotiation;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {

```



```

failure-action (clear-binding | clear-binding-if-interface-up | log-only);
method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
  layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
  }
}
}
overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-negotiation-match incoming-interface;
  delegated-pool;
  delete-binding-on-renegotiation;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
}
}
reauthenticate (<lease-renewal> <remote-id-mismatch>);
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  support-option-pd-exclude;
  timeout timeout-value;
  token token-value;

```

```
    trigger {
        radius-disconnect;
    }
}
requested-ip-network-match subnet-mask;
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
[edit logical-systems *logical-system-name* system services [dhcp-local-server](#)],
[edit routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
[edit system services [dhcp-local-server](#)]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Configure DHCPv6 local server options on the router or switch to enable the router or switch to function as a server for the DHCP protocol for IPv6. The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *DHCPv6 Local Server Overview*

dhcpv6 (DHCP Relay Agent)

```
Syntax  dhcpv6 {
    access-profile profile-name;
    active-server-group server-group-name;
}
authentication {
    password password-string;
    username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description (device-interface | logical-interface);
        interface-name interface-name;
        logical-system-name;
        mac-address mac-address;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
    }
}
bulk-leasquery {
    attempts number-of-attempts;
    timeout seconds;
    trigger automatic;
}
duplicate-clients incoming-interface;
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name;
```

```
    logical-system-name;
    mac-address;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
  }
}
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
interface interface-name {
  access-profile profile-name;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  exclude;
  overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
}
lease-time-validation {
  lease-time-threshold seconds;
  violation-action action;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
```

```

    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
}
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
    }
}

```

```

        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
}
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
        route-suppression;
        service-profile dynamic-profile-name;
    }
}
no-snoop;
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
}

```

```

interface-client-limit number;
no-allow-snooped-clients;
no-bind-on-request;
relay-source interface-name;
send-release-on-delete;
}
relay-agent-interface-id {
include-irb-and-l2;
keep-incoming-interface-id ;
no-vlan-interface-name;
prefix prefix;
use-interface-description (logical | device);
use-option-82;
}
relay-agent-remote-id {
include-irb-and-l2;
keep-incoming-interface-id ;
no-vlan-interface-name;
prefix prefix;
use-interface-description (logical | device);
use-option-82 <strict>;
}
relay-option {
option-number option-number;
default-action {
drop;
forward-only;
relay-server-group relay-server-group;
}
equals (ascii ascii-string | hexadecimal hexadecimal-string) {
drop;
forward-only;
relay-server-group relay-server-group;
}
starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
drop;
forward-only;
relay-server-group relay-server-group;
}
}
relay-option-vendor-specific{
host-name;
location;
remote-id-mismatch disconnect;
route-suppression;
server-group {
server-group-name {
server-ip-address;
}
}
server-response-time seconds;
service-profile dynamic-profile-name;
}

```

Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3 for EX Series switches. Support for forward-snooped-clients introduced in Junos OS Release 15.1X53-D56 for EX Series switches and Junos OS Release 17.1R1.</p>
Description	<p>Configure DHCPv6 relay options on the router or switch and enable the router or switch to function as a DHCPv6 relay agent. A DHCPv6 relay agent forwards DHCPv6 request and reply packets between a DHCPv6 client and a DHCPv6 server.</p> <p>The DHCPv6 relay agent server is fully compatible with the extended DHCP local server and DHCP relay agent. However, the options configured with the dhcpv6 statement are incompatible with the DHCP/BOOTP relay agent options configured with the bootp statement. As a result, the DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot be enabled on the router or switch at the same time.</p> <p>The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 431• <i>DHCPv6 Relay Agent Overview</i>• <i>Using External AAA Authentication Services with DHCP</i>

dial-options

Syntax	<pre>dial-options { ipsec-interface-id <i>name</i>; l2tp-interface-id <i>name</i>; (shared dedicated); }</pre>
Hierarchy Level	<p>[edit interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit interfaces <i>si-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>si-fpc/pic/port</i> unit <i>logical-unit-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. The [edit ...si-...] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.
Options	<p>dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.</p> <p>ipsec-interface-id <i>name</i>—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the [edit access profile <i>name</i> client * ike] hierarchy level.</p> <p>l2tp-interface-id <i>name</i>—Interface identifier that must be replicated at the [edit access profile <i>name</i>] hierarchy level.</p> <p>shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Identifier for Logical Interfaces that Provide L2TP Services</i> • <i>Configuring Dynamic Endpoints for IPsec Tunnels</i> • Configuring Options for the LNS Inline Services Logical Interface on page 253

dial-options (Dynamic Profiles)

Syntax	<pre>dial-options { ipsec-interface-id <i>name</i>; l2tp-interface-id <i>name</i>; (shared dedicated); }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the options for configuring logical interfaces in dynamic profiles for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.
Options	<p>dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.</p> <p>ipsec-interface-id <i>name</i>—Interface identifier for group of dynamic peers. This identifier must be replicated at the [edit access profile <i>name</i> client * <i>ike</i>] hierarchy level. This option is not currently supported for dynamic profiles.</p> <p>l2tp-interface-id <i>name</i>—(MX Series routers only) L2TP interface identifier that must be replicated at the [edit access profile <i>name</i>] hierarchy level.</p> <p>shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Dynamic Profile for Dynamic LNS Sessions on page 285

disable-calling-number-avp (L2TP LAC)

Syntax	disable-calling-number-avp;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Prevent the LAC from sending L2TP Calling Number AVP 22 in incoming-call request (ICRQ) packets to the LNS. By default, the LAC in an L2TP network generates this AVP from the Calling-Station-Id and sends it to the LNS.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Preventing the LAC from Sending Calling Number AVP 22 to the LNS on page 235

disable-failover-protocol (L2TP)

Syntax disable-failover-protocol;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 11.2.
Statement deprecated in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1.

Description Configure the LAC or LNS to use only the silent failover method when resynchronizing with its peer in the event of a control plane failover. This statement prevents the default behavior, where the LAC first attempts to negotiate the failover protocol when it establishes control connections with the peer. If the remote peer does not support the failover protocol, then the LAC falls back on the silent failover method. Including this configuration is useful when the peers configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it.



BEST PRACTICE: We recommend that you include this statement on both the LAC and LNS to prevent the use of failover protocol. When failover protocol is used, the nonfailed peer (LAC or LNS) keeps the tunnel open with the failed peer, in case the failed peer is able to recover from the failure and resynchronize with the nonfailed peer. This behavior keeps the tunnel up and the subscribers logged in while traffic is not flowing, preventing service level agreements from being met.

Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the **disable-failover-protocol** statement is deprecated and no longer needs to be used. The default failover resynchronization method is changed to silent failover, rather than the previous default method of failover-protocol-fall-back-to-silent-failover. The new default method conforms to our recommendation to use silent failover. Consequently, there is no need to disable the failover protocol. Configurations that include this statement are still supported when you upgrade to a release in which it is deprecated; The CLI informs you of the deprecation if the statement is included.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release History Table

Release	Description
15.1R6	Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the disable-failover-protocol statement is deprecated and no longer needs to be used.

- Related Documentation**
- [Configuring the L2TP Peer Resynchronization Method on page 295](#)

drain

Syntax	drain;
Hierarchy Level	<p>[edit services l2tp], [edit services l2tp destination address <i>ip-address</i>], [edit services l2tp destination address <i>ip-address</i> routing-instance <i>routing-instance-name</i>], [edit services l2tp destination name <i>destination-name</i>], [edit services l2tp tunnel name <i>name</i>], [edit services l2tp tunnel name <i>name</i> address <i>ip-address</i>], [edit services l2tp tunnel name <i>name</i> address <i>ip-address</i> routing-instance <i>routing-instance-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Prevent the creation of new sessions, destinations, and tunnels globally at an L2TP access concentrator (LAC) or an L2TP network server (LNS). Prevent the creation of new tunnels and sessions for a specific destination. Prevent the creation of new sessions for a specific tunnel.
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring L2TP Drain on page 166 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

dual-stack-group (DHCP Local Server)

```
Syntax  dual-stack-group name {
        access-profile access-profile;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        on-demand-address-allocation;
        protocol-master (inet | inet6);
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        service-profile service-profile;
    }
```

Hierarchy Level [edit logical-systems *name* routing-instances *name* system services [dhcp-local-server](#)],
 [edit logical-systems *name* system services [dhcp-local-server](#)],
 [edit routing-instances *name* system services [dhcp-local-server](#)],
 [edit system services [dhcp-local-server](#)]

Release Information Statement introduced in Junos OS Release 17.3R1 on MX Series routers.

Description	<p>Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP local server dual-stack, and names the dual-stack group.</p> <p>When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.</p>
Options	<p><i>name</i>—Name of the dual-stack group.</p> <p>The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement in the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Single-Session DHCP Local Server Dual-Stack Overview</i>• DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98• <i>Configuring Reauthentication for DHCP Local Server</i>• <i>RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers</i>

dual-stack-group (DHCP Relay Agent)

```
Syntax  dual-stack-group name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        protocol-master;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
```



```

        use-option-82 <strict>;
    }
    service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit forwarding-options [dhcp-relay](#)],
 [edit logical-systems *logical-system-name* forwarding-options [dhcp-relay](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options [dhcp-relay](#)],
 [edit routing-instances *routing-instance-name* forwarding-options [dhcp-relay](#)]

Release Information Statement introduced in Junos OS Release 15.1.

Description Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP dual stack, and names the dual stack group.

The group is assigned to each leg of the DHCP dual-stack with the *dual-stack* statement in the [overrides](#) stanza. When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

Options *name*—Name of the dual-stack group.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement in the configuration.

Related Documentation

- *Single-Session DHCP Dual-Stack Overview*
- *Configuring Single-Session DHCP Dual-Stack Support*
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)

duplicate-clients (DHCPv6 Local Server and Relay Agent)

Syntax	duplicate-clients incoming-interface;
Hierarchy Level	[edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ...], [edit routing-instances <i>routing-instance-name</i> ...], [edit system services dhcp-local-server dhcpv6]
Release Information	Statement introduced in Junos OS Release 16.1.
Description	<p>Specify the criteria that the jdhcpd process uses to support duplicate clients. The router uses the additional criteria to distinguish between the duplicate clients.</p> <p>Duplicate clients have the same DUID (DHCP unique identifier). Typically, the router treats a request from a duplicate client as a renegotiation, and replaces the existing client entry with a new entry. However, in some cases, the duplicate request is from a different client, and replacement is not desired. When you enable duplicate client support, the router uses the additional criteria to distinguish between the two clients, and grants a lease to the new client while retaining the original client entry.</p>



NOTE: The only supported differentiating criterion is **incoming-interface**.



BEST PRACTICE: To allow duplicate clients over the incoming interface for DHCPv6 relay, you must configure the **relay-agent-interface-id** statement to cause the DHCP relay agent to insert the DHCPv6 Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.

Do not configure the **use-interface-description** statement, because option 18 must include the interface name rather than an interface description.



CAUTION: We recommend that you do not enable or disable duplicate client support mode when clients are bound, because different client keys are used to store the clients in the database depending on the mode. Changing the mode removes clients from the database and then adds them back with a different key.

Additionally, disabling duplicate client support mode causes all duplicate clients to be deleted.

Options	incoming-interface —Allow duplicate clients when the incoming DHCPv6 requests are received over different underlying interfaces.
Required Privilege	system—To view this statement in the configuration.
Level	system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• DHCPv6 Duplicate Client DUIDs on page 67• Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs on page 68

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)

Syntax duplicate-clients-in-subnet (incoming-interface | option-82);

Hierarchy Level [edit forwarding-options [dhcp-relay](#)],
 [edit logical-systems *logical-system-name* forwarding-options [dhcp-relay](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options [dhcp-relay](#)],
 [edit routing-instances *routing-instance-name* forwarding-options [dhcp-relay](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
 [edit logical-systems *logical-system-name* system services [dhcp-local-server](#)],
 [edit routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
 [edit system services [dhcp-local-server](#)]

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure how the router distinguishes between duplicate clients in the same subnet. Duplicate clients are defined as clients that have the same hardware address or client ID.



NOTE: You must configure the duplicate-clients-in-subnet statement identically for both the DHCP local server ([edit forwarding-options [dhcp-relay](#)]) and the DHCP relay agent ([edit system services [dhcp-local-server](#)]).

Options **incoming-interface**—Use the incoming interface information in packets to differentiate between duplicate clients.

option-82—Use the option 82 information to differentiate between duplicate clients. Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

- Related Documentation**
- [DHCPv4 Duplicate Client In Subnet Overview on page 61](#)
 - [Guidelines for Configuring Support for DHCPv4 Duplicate Clients on page 62](#)
 - [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces on page 64](#)
 - [Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information on page 63](#)

dynamic-profile (L2TP)

Syntax	dynamic-profile <i>profile-name</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Assign a dynamic profile to the tunnel group for dynamic LNS sessions.
Options	<i>profile-name</i> —Name of the dynamic profile for the tunnel group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Dynamic Profile for Dynamic LNS Sessions on page 285


dynamic-profile (PPP)

Syntax	dynamic-profile <i>profile-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]
Release Information	Statement introduced in Junos OS Release 9.5. Support for MLPPP on LSQ interfaces introduced in Junos OS Release 10.2.
Description	Specify the dynamic profile that is attached to the interface. On the MX Series routers, this statement is supported on PPPoE interfaces only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Dynamic Profiles Overview</i>• <i>Configuring a Basic Dynamic Profile</i>• Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121• <i>Attaching Dynamic Profiles to MLPPP Bundles</i>• For hardware requirements, see <i>Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces</i>

enable-ipv6-services-for-lac (L2TP)

Syntax	enable-ipv6-services-for-lac;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 17.1.
Description	Enable the LAC to create the IPv6 address family (inet6) when establishing a tunnel for subscribers, allowing IPv6 filters to be applied. By default, the LAC requires only family inet to enable forwarding into an IP tunnel. It can apply IPv4 firewall filters to the session. Even when family inet6 is included in the dynamic profile, by default it is not created and IPv6 firewall filters cannot be applied.
Default	Disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling the LAC for IPv6 Services on page 212• Configuring an L2TP LAC on page 173

enable-snmp-tunnel-statistics (L2TP)

Syntax	enable-snmp-tunnel-statistics;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 12.1R4 and supported in later 12.1Rx releases. Statement supported in Junos OS Release 12.2R2 and later 12.2Rx releases. (Not supported in Junos OS Release 12.2R1.) Statement supported in Junos OS Release 12.3 and later releases.
Description	Enable collection of L2TP tunnel and global counters for SNMP statistics. <div> NOTE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.</div>
Default	Disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Tunnel and Global Counters for SNMP Statistics Collection on page 303

encapsulation (Logical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-ccc-vc-mux atm-cisco-nlpid atm-mlppp-llc atm-nlpid atm-ppp-llc atm-ppp-vc-mux atm-snap atm-tcc-snap atm-tcc-vc-mux atm-vc-mux ether-over-atm-llc ether-vpls-over-atm-llc ether-vpls-over-fr ether-vpls-over-ppp ethernet ethernet-ccc ethernet-vpls ethernet-vpls-fr frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-ppp frame-relay-tcc gre-fragmentation multilink-frame-relay-end-to-end multilink-ppp ppp-over-ether ppp-over-ether-over-atm-llc vlan-bridge vlan-ccc vlan-vci-ccc vlan-tcc vlan-vpls vxlan);
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>], [edit interfaces rlsq <i>number</i> unit <i>logical-unit-number</i>] [edit protocols evpn]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (ethernet , vlan-ccc , and vlan-tcc options only). Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers. Only the atm-ccc-cell-relay and atm-ccc-vc-mux options are supported on ACX Series routers. Statement introduced in Junos OS Release 17.3R1 for QFX10000 Series switches (ethernet-ccc and vlan-ccc options only).
Description	Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the ccc family only.</p> <p>atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.</p> <p>atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.</p> <p>atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.</p>

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.



NOTE: The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

vxlan—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
---------------------------------	---

Related Documentation

- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring the Encapsulation for Layer 2 Switching TCCs*
- *Configuring Interface Encapsulation on Logical Interfaces*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Circuit and Translational Cross-Connects Overview*
- *Identifying the Access Concentrator*
- *Configuring ATM Interface Encapsulation*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM for Subscriber Access*
- *Understanding CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*


enforce-strict-scale-limit-license (Subscriber Management)

Syntax	enforce-strict-scale-limit-license;
Hierarchy Level	[edit system services subscriber-management]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Configure the router to strictly enforce the subscriber scaling license, and to not allow the normal grace period. No additional subscribers are allowed to log in after the number of subscribers reaches the maximum allowed for the license.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Router to Strictly Enforce the Subscriber Scaling License</i>

equals (Dynamic Profile)

Syntax	equals <i>expression</i> ;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> variables <i>variable-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Configure an expression for a user-defined variable that is evaluated at run time and returned as the variable value.
Options	<i>expression</i> —Expression evaluated to return a value for the user-defined variable.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>User-Defined Variables</i> • <i>Using Variable Expressions in User-Defined Variables</i> • <i>Configuring User-Defined Dynamic Variables in Dynamic Profiles</i>

failover-resync

Syntax	failover-resync (failover-protocol silent-failover)
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 17.2R1 on MX Series.
Description	<p>Configure the method used by the LAC or the LNS to resynchronize with its peer in the event of a control plane failover. Failure can be the result of a Routing Engine switchover, a daemon restart, or some other cause. During tunnel setup, the L2TP endpoints negotiate the resynchronization method; silent failover is the default.</p> <p>With the silent failover method, only the failed endpoint is involved in recovering the tunnels and sessions; the nonfailed endpoint remains unaware of the failure.</p> <p>With the failover protocol method, the nonfailed endpoint keeps the tunnel open with the failed peer, in case the failed peer is able to recover from the failure and resynchronize with the nonfailed peer. The detection of tunnel keepalive failures is delayed. This behavior keeps the tunnel up and the subscribers logged in while traffic is not flowing, preventing service level agreements from being met.</p>
	<div> BEST PRACTICE: Use the default method, silent failover.</div>
	<p>This statement supersedes the deprecated statement, disable-failover-protocol.</p>
Options	<p>failover-protocol—Specify the L2TP failover protocol as the resynchronization method, but fall back to silent failover if the other endpoint does not support it.</p> <p>silent-failover—Specify silent failover as the resynchronization method.</p> <p>Default: silent-failover</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the L2TP Peer Resynchronization Method on page 295• L2TP Failover and Peer Resynchronization on page 294

failover-within-preference (L2TP LAC)

Syntax	failover-within-preference;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Enable L2TP LAC tunnel selection within a preference level. When the router is unable to connect to a destination at a given preference level, it attempts to connect to another destination at the same level. By default, when a connection attempt fails at one preference level, the next attempt is made at the next lower level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LAC Tunnel Selection Failover Within a Preference Level on page 210• Configuring the L2TP LAC Tunnel Selection Parameters on page 209

failure-action

Syntax	failure-action (clear-binding clear-binding-if-interface-up log-only);
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection], [edit system services dhcp-local-server dhcpv6 liveness-detection], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the action the router (or switch) takes when a liveness detection failure occurs.
Options	<p>Default: clear-binding</p> <p>clear-binding—The DHCP client session is cleared when a liveness detection failure occurs, except when maintain-subscribers interface-delete setting is configured and active.</p> <p>clear-binding-if-interface-up—The DHCP client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up. Use this setting to distinguish failures from between a liveness detection failure due to a local network error, and a host disconnecting from the network. If the client binding is in the maintain-binding Finite State Machine (FSM) state when the liveness detection failure detection occurs, then the binding is not deleted. Not supported for Layer 2 ARP/ND liveness detection on MX Series routers.</p> <p>log-only—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure and maintain the client binding. Not supported for Layer 2 ARP/ND liveness detection on MX Series routers.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Overview on page 85 • Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92 • Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87 • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

flexible-vlan-tagging

Syntax	flexible-vlan-tagging;
Hierarchy Level	[edit interfaces <i>aex</i>], [edit interfaces <i>ge-fpc/pic/port</i>], [edit interfaces <i>et-fpc/pic/port</i>], [edit interfaces <i>ps0</i>], [edit interfaces <i>xe-fpc/pic/port</i>]
Release Information	Statement introduced in Junos OS Release 8.1. Support for aggregated Ethernet added in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
Description	<p>Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.</p> <p>This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.</p> <p>This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling VLAN Tagging</i> • <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i> • <i>Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces</i>

forward-snooped-clients (DHCP Local Server)

Syntax	forward-snooped-clients (all-interfaces configured-interfaces non-configured-interfaces);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit system services dhcp-local-server]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure how the DHCP local server filters and handles DHCP snooped packets on the specified interfaces.
Options	all-interfaces —Perform the action on all interfaces. configured-interfaces —Perform the action only on interfaces that are configured as part of an interface group. non-configured-interfaces —Perform the action only on interfaces that are not configured as part of a group.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• DHCP Snooping Support on page 41• Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 43

forward-snooped-clients (DHCP Relay Agent)

Syntax	<code>forward-snooped-clients (all-interfaces configured-interfaces non-configured-interfaces);</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level introduced in Junos OS Release 15.1X53-D56 for EX Series switches and Junos OS Release 17.1R1.</p>
Description	<p>Configure how DHCP relay agent filters and handles DHCP snooped packets on the specified interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the forward-snooped-clients configuration and the configuration of either the allow-snooped-clients statement or the no-allow-snooped-clients statement.</p> <p>The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on non-configured interfaces.</p>
Options	<p>all-interfaces—Perform the action on all interfaces.</p> <p>Default: On EX Series switches, the action is performed on all interfaces by default.</p> <p>configured-interfaces—Perform the action only on interfaces that are configured as part of an interface group.</p> <p>non-configured-interfaces—Perform the action only on interfaces that are not part of a group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Snooping Support on page 41 • Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48

fpc (MX Series 3D Universal Edge Routers)

```
Syntax  fpc slot-number {
        inline-services {
            flow-table-size {
                ipv4-flow-table-size units;
                ipv4-flow-table-size units;
                ipv6-extended-attrib;
            }
        }
        ir-mode (R | IR);
        pic number {
            inline-services {
                bandwidth (1g | 10g);
            }
            port-mirror-instance port-mirroring-instance-name-pic-level;
            tunnel-services {
                bandwidth (1g | 10g)
            }
        }
        port-mirror-instance port-mirroring-instance-name-fpc-level;
    }
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 8.2.
port-mirror-instance option added in Junos OS Release 9.3.
ipv6-extended-attrib option added in Junos OS Release 14.2 for MX Series routers.

Description Configure properties for the DPC or MPC and corresponding Packet Forwarding Engines to create tunnel interfaces.

(MX Series Virtual Chassis only) When you configure chassis properties for MPCs installed in a Virtual Chassis member router, statements included at the **[edit chassis member member-id fpc slot slot-number]** hierarchy level apply to the MPC in the specified slot number only on the specified member router in the Virtual Chassis. Statements included at the **[edit chassis fpc slot slot-number]** hierarchy level apply to the MPCs in the specified slot number on *each* member router in the Virtual Chassis.



BEST PRACTICE: To ensure that the statement you use to configure MPC chassis properties in an MX Series Virtual Chassis applies to the intended member router and MPC, we recommend that you always include the **member member-id** option before the **fpc** statement, where **member-id** is 0 or 1 for a two-member MX Series Virtual Chassis.

Options **fpc slot-number**—Specify the slot number of the DPC.
Range: 0 through 11

pic *number*—Specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

Range: 0 through 4

port-mirror-instance *port-mirroring-instance-name-fpc-level*—Associate a port-mirroring instance with the DPC and its corresponding PICs. The port-mirroring instance is configured under the **[edit forwarding-options port-mirroring]** hierarchy level.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Port-Mirroring Instances on MX Series 3D Universal Edge Routers*
- [Enabling Inline Service Interfaces on page 251](#)
- *Virtual Chassis Components Overview*

gateway-name (LNS Local Gateway)

Syntax `gateway-name gateway-name;`

Hierarchy Level [edit services l2tp tunnel-group *group-name* [local-gateway](#)]

Release Information Statement introduced in Junos OS Release 12.2.

Description Specify the gateway name for the LNS, which the LNS returns to the LAC in response to the LAC's SCCRQ message. This name must match the remote gateway name configured on the LAC, or the tunnel cannot be established.

Options ***gateway-name***—Name of the LNS.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)

gateway-name (Tunnel Profile Remote Gateway)

Syntax	<code>gateway-name <i>server-name</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i> remote-gateway]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the hostname expected by the remote gateway—the LNS—from the source gateway—the LAC—when you set up a tunnel.
Options	<i>server-name</i> —Name of the LNS.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

gateway-name (Tunnel Profile Source Gateway)

Syntax	<code>gateway-name <i>client-name</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i> source-gateway]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the hostname provided by the source gateway—the LAC—to the remote gateway—the LNS—when you set up a tunnel.
Options	<i>client-name</i> —Name of the LAC.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

gres-route-flush-delay (Subscriber Management)

Syntax	gres-route-flush-delay;
Hierarchy Level	[edit system services subscriber-management]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, configure the router to wait 180 seconds (3 minutes) before removing (flushing) static or dynamic access routes and access-internal routes from the forwarding table after a graceful Routing Engine switchover (GRES) has taken place.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 108• Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 107

group (DHCP Local Server)

```
Syntax  group group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description (device-interface | logical-interface);
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                relay-agent-interface-id
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
            primary-profile-name>;
        interface interface-name {
            access-profile profile-name;
            exclude;
            overrides {
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-discover-match <option60-and-option82>;
                client-negotiation-match incoming-interface;
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                    }
                }
            }
        }
    }
```



```

        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
}
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82>;
    client-negotiation-match incoming-interface;
    delegated-pool;
    delete-binding-on-renegotiation;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
}
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit system services [dhcp-local-server](#)],
[edit system services [dhcp-local-server dhcpv6](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system
services [dhcp-local-server](#) ...],
[edit logical-systems *logical-system-name* system services [dhcp-local-server](#) ...],
[edit routing-instances *routing-instance-name* system services [dhcp-local-server](#) ...]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Options *group-name*—Name of the group.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege system—To view this statement in the configuration.
Level system-control—To add this statement to the configuration.

Related Documentation

- *Extended DHCP Local Server Overview*
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)

group (DHCP Relay Agent)

```
Syntax  group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
    }
```

```

}
overrides {
    allow-no-end-option;
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holdddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;

```

```

always-write-option-82;
asymmetric-lease-time seconds;
asymmetric-prefix-lease-time seconds;
client-discover-match <option60-and-option82>;
client-negotiation-match incoming-interface;
disable-relay;
dual-stack dual-stack-group-name;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-bind-on-request;
proxy-mode;
relay-source
replace-ip-source-with;
send-release-on-delete;
trust-option-82;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    remote-id {
        prefix prefix;

```

```

        use-interface-description (logical | device);
    }
    server-id-override
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level	<p>[edit forwarding-options dhcp-relay],</p> <p>[edit forwarding-options dhcp-relay dhcpv6],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Options	<p>group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.</p> <p>The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Extended DHCP Relay Agent Overview</i> • <i>Understanding the Extended DHCP Relay Agent for EX Series Switches</i> • <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i> • <i>Configuring Group-Specific DHCP Relay Options</i> • <i>Grouping Interfaces with Common DHCP Configurations</i> • <i>Using External AAA Authentication Services with DHCP</i> • <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i> • DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98

group-profile (Group Profile)

Syntax `group-profile profile-name {`
 `l2tp {`
 `interface-id interface-id;`
 `lcp-renegotiation;`
 `local-chap;`
 `maximum-sessions-per-tunnel number;`
 `}`
 `ppp {`
 `cell-overhead;`
 `encapsulation-overhead bytes;`
 `framed-pool pool-id;`
 `idle-timeout seconds;`
 `interface-id interface-id;`
 `keepalive seconds;`
 `ppp-options {`
 `aaa-options aaa-options-name;`
 `chap;`
 `ignore-magic-number-mismatch;`
 `initiate-ncp (ip | ipv6 | dual-stack-passive)`
 `ipcp-suggest-dns-option;`
 `mru;`
 `mtu;`
 `pap;`
 `peer-ip-address-optional;`
 `}`
 `primary-dns primary-dns;`
 `primary-wins primary-wins;`
 `secondary-dns secondary-dns;`
 `secondary-wins secondary-wins;`
 `}`
`}`

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the group profile.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *profile-name*—Name assigned to the group profile.

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	<ul style="list-style-type: none">• <i>Configuring the Group Profile for Defining L2TP Attributes</i>• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

Syntax	<pre> hierarchical-scheduler { implicit-hierarchy; maximum-hierarchy-levels <i>number</i>; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 10.1.</p> <p>implicit-hierarchy option added in Junos OS Release 13.1.</p> <p>Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.</p> <p>Support for up to four hierarchy levels added in Junos OS Release 16.1.</p>
Description	<p>Configure hierarchical scheduling options on the interface.</p> <p>The statement is supported on the following interfaces:</p> <ul style="list-style-type: none"> • MIC and MPC interfaces in MX Series routers • GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers <p>To enable hierarchical scheduling on MX Series routers, configure the hierarchical-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.</p>
Options	<p>implicit-hierarchy—Configure four-level hierarchical scheduling. When you include the implicit-hierarchy option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, level 3, and level 4. The implicit-hierarchy option is supported only on MPC/MIC subscriber interfaces and interface sets on MX Series routers.</p> <p>maximum-hierarchy-levels <i>number</i>—Specify the maximum number of hierarchical scheduling levels allowed for node scaling, from 2 through 4 levels. The default number of levels is 3. The maximum-hierarchy-levels option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets on MX Series routers.</p> <ul style="list-style-type: none"> • If you set maximum-hierarchy-levels to 2, interface sets are not allowed. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors. • If you do not include the maximum-hierarchy-levels option, keeping the default number of hierarchy levels at 3, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.

Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Hierarchical CoS for Subscriber Interfaces• Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links• Configuring Hierarchical Schedulers for CoS• Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface• Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 323

holddown-interval

Syntax	holddown-interval <i>milliseconds</i> ;
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.
Options	<i>milliseconds</i> —Interval specifying how long a BFD session must remain up before a state change notification is sent. Range: 0 through 255,000 Default: 0
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

hello-interval (L2TP)

Syntax	<code>hello-interval <i>seconds</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the keepalive timer for L2TP tunnels.
Options	<p><i>seconds</i>—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of 0 means that no hello messages are sent.</p> <p>Range: 0 through 3600</p> <p>Default: 60 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Timers for L2TP Tunnels • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274

identification (Tunnel Profile)

Syntax	<code>identification <i>name</i>;</code>
Hierarchy Level	<code>[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the assignment ID of an L2TP tunnel. L2TP sessions with the same tunnel assignment identification and destination are grouped into the same tunnel.
Options	<i>name</i> —Tunnel assignment ID; string of up to 32 alphanumeric characters.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Tunnel Profile for Subscriber Access on page 207

idle-timeout (Access)


Syntax	<code>idle-timeout seconds;</code>
Hierarchy Level	<code>[edit access group-profile <i>profile-name</i> ppp],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> ppp]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons: <ul style="list-style-type: none">• There is no ingress traffic on the PPP session.• There is no egress traffic.• There is neither ingress or egress traffic on the PPP session.• There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.
Options	seconds —Number of seconds a user can remain idle before the session is terminated. Range: 0 through 4,294,967,295 seconds Default: 0



NOTE: The `[edit access]` hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Group Profile for Defining L2TP Attributes• Configuring PPP Properties for a Client-Specific Profile• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245

idle-timeout (L2TP)

Syntax	<code>idle-timeout seconds;</code>
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Specify how long a tunnel is active after its last session is terminated. The timer starts when the session is terminated and the tunnel is disconnected when the timer expires.
	<div>  <p>BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing <code>no services l2tp tunnel idle-timeout</code>.</p> </div>
Options	<p><i>seconds</i>—Length of the idle timeout. A value of 0 creates a persistent tunnel; that is, the tunnel remains active indefinitely until the remote peer disconnects it or you issue the clear services l2tp tunnel command.</p> <p>Range: 0 through 86,400</p> <p>Default: 60</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Setting the L2TP Tunnel Idle Timeout on page 163 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

ignore-magic-number-mismatch (Access Group Profile)

Syntax ignore-magic-number-mismatch;

Hierarchy Level [edit access **group-profile name** **ppp ppp-options**]

Release Information Statement introduced in Junos OS Release 18.1R1 on MX Series routers.

Description Prevent the Packet Forwarding Engine from performing a validation check for magic numbers received in LCP keepalive (Echo-Request/Echo-Reply) exchanges for a group of tunneled PPP subscribers at the LNS.
A mismatch occurs when the PPP magic number received from a remote peer in the keepalive exchange does not match the value agreed upon during LCP negotiation. Disabling the validation check prevents PPP from terminating the session when an unexpected number is received. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.



NOTE: Because magic number validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends the local peer's magic number, which would indicate a loopback or other network issue. This is considered to be an unlikely situation, because LCP negotiation completed successfully, meaning no loopback was present at that time.



NOTE: You can also configure this behavior in a dynamic PPP profile. When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

This means that **ignore-magic-number-mismatch** configured in a group profile is not applied when the dynamic profile includes any PPP option, even when the dynamic profile does not include **ignore-magic-number-mismatch** statement.



NOTE: This statement is not supported on static interfaces.

Required Privilege access—To view this statement in the configuration.
Level access-control—To add this statement to the configuration.

Related Documentation

- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)
- For dynamic profiles: [Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests on page 116](#)

ignore-magic-number-mismatch (Dynamic Profiles)

Syntax	ignore-magic-number-mismatch;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]
Release Information	Statement introduced in Junos OS Release 18.1R1 on MX Series routers.
Description	<p>Prevent the Packet Forwarding Engine from performing a validation check for magic numbers received in LCP keepalive (Echo-Request/Echo-Reply) exchanges for dynamic PPP subscriber connections terminated at the router or for dynamic tunneled PPP subscribers on LNS inline service interfaces.</p> <p>A mismatch occurs when the PPP magic number received from a remote peer in the keepalive exchange does not match the value agreed upon during LCP negotiation. Disabling the validation check prevents PPP from terminating the session when an unexpected number is received. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.</p> <p>For dynamic PPP subscriber connections terminated at the router, configure the statement at the [edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options] hierarchy level.</p> <p>For dynamic tunneled PPP subscribers on LNS inline service interfaces, configure the statement at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options] hierarchy level.</p>



NOTE: Because magic number validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends the local peer's magic number, which would indicate a loopback or other network issue. This is considered to be an unlikely situation, because LCP negotiation completed successfully, meaning no loopback was present at that time.



NOTE: You can also configure this behavior in an L2TP group profile that applies to tunneled PPP subscribers at the LNS. When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

This means that `ignore-magic-number-mismatch` configured in a group profile is not applied when the dynamic profile includes any PPP option, even when the dynamic profile does not include `ignore-magic-number-mismatch`.



NOTE: This statement is not supported on static interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges on page 120](#)
- [Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests on page 116](#)
- For L2TP group profiles: [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)

initiate-ncp (Dynamic and Static PPP)

Syntax	<code>initiate-ncp (ip ipv6 dual-stack-passive);</code>
Hierarchy Level	<code>[edit access group-profile <i>profile-name</i> ppp ppp-options],</code> <code>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options],</code> <code>[edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit</code> <code> "\$junos-interface-unit" ppp-options],</code> <code>[edit interfaces pp0 unit <i>logical-unit-number</i> ppp-options],</code> <code>[edit interfaces <i>si-fpc/pic/port</i> unit <i>logical-unit-number</i> ppp-options]</code>
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Configure PPP Network Control Protocol (NCP) negotiation mode (active or passive) for dynamic and static IPv4 and IPv6 PPP subscriber interfaces. You can also configure PPP NCP negotiation mode for the PPP server in an IPv4/IPv6 dual-stack configuration.
Options	<p>dual-stack-passive—Enable passive PPP NCP negotiation for the PPP server in an IPv4/IPv6 dual-stack configuration. The initiate-ncp dual-stack-passive statement overrides the initiate-ncp ip and initiate-ncp ipv6 statements if they are configured in an IPv4/IPv6 dual-stack configuration.</p> <p>ip—Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv4 (inet) protocol address family, and for which IPv4 address attributes are assigned during authorization. By default, dynamic and static IPv4 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the initiate-ncp ip statement to enable active PPP NCP negotiation for the IPv4 subscriber interface.</p> <p>ipv6—Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv6 (inet6) protocol address family, and for which IPv6 address attributes are assigned during authorization. By default, dynamic and static IPv6 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the initiate-ncp ipv6 statement to enable active PPP NCP negotiation for the IPv6 subscriber interface.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the PPP Network Control Protocol Negotiation Mode on page 134• PPP Network Control Protocol Negotiation Mode Overview on page 129

inline-services (FPC Level)

Syntax

```
inline-services {
  flow-table-size {
    ipv4-flow-table-size units;
    ipv6-extended-attrib;
    ipv6-flow-table-size units;
    mpls-flow-table-size units;
    vpls-flow-table-size units;
  }
}
```

Hierarchy Level [edit chassis **fpc** *slot-number*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Enable inline services on MPCs, configured at the FPC level. To enable inline services that are specified at the PIC level, see the configuration statement [inline-services \(PIC level\)](#).



NOTE: On MX80 routers and MX Series routers with Trio-based FPCs, when ingress queuing is enabled for a PIC, tunnel services and inline services are not supported on the same PIC.

Options The remaining statements are defined separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Enabling Inline Service Interfaces on page 251](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#)

inline-services (PIC level)

Syntax	<code>inline-services { bandwidth (1g 10g 20g 30g 40g); }</code>
Hierarchy Level	[edit chassis fpc <i>slot-number</i> pic <i>number</i>]
Release Information	Statement introduced in Junos OS Release 11.4. 20g , 30g , and 40g options added in Junos OS Release 14.1R3.
Description	Enable inline services on PICs residing on MPCs. To enable inline services that are specified at the fpc level, see configuration statement inline-services (FPC Level) The remaining statement is explained separately. See CLI Explorer .
Options	The option is 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">• Enabling Inline Service Interfaces on page 251• Configuring an L2TP LNS with Inline Service Interfaces on page 239

input-hierarchical-policer

Syntax	<code>input-hierarchical-policer <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> layer2-policer], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer],
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a hierarchical policer to the Layer 2 input traffic for all protocol families at the physical or logical interface.
Options	<i>policer-name</i> —Name of the hierarchical policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Hierarchical Policers</i> • <i>layer2-policer (Hierarchical Policar)</i>

interface (Dynamic Routing Instances)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Assign the specified interface to the dynamically created routing instance.
Options	<i>interface-name</i> —The interface name variable (<i>\$junos-interface-name</i>). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	

interface (L2TP Service Interfaces)

Syntax	<code>interface <i>service-interface-name</i>;</code>
Hierarchy Level	[edit services service-device-pools pool <i>pool-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify a service interface assigned to a service interface pool. You specify more than one interface for each pool; the interfaces are used by an L2TP tunnel group to balance traffic loads.
Options	<i>service-interface-name</i> —Name of the service interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 284• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274

interface-id

Syntax	<code>interface-id <i>interface-id</i>;</code>
Hierarchy Level	[edit access group-profile <i>profile-name</i> l2tp], [edit access group-profile <i>profile-name</i> ppp], [edit access profile <i>profile-name</i> client <i>client-name</i> ike], [edit access profile <i>profile-name</i> client <i>client-name</i> l2tp], [edit access profile <i>profile-name</i> client <i>client-name</i> ppp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the interface identifier.
Options	<i>interface-id</i> —Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the [edit interfaces <i>interface-name</i> unit <i>local-unit-number</i> dial-options] hierarchy level. For more information about the interface ID, see <i>Services Interface Naming Overview</i> .
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring L2TP for a Group Profile</i> • <i>Configuring the Group Profile for Defining L2TP Attributes</i> • <i>Configuring L2TP Properties for a Client-Specific Profile</i> • <i>Configuring PPP Properties for a Client-Specific Profile</i> • <i>Configuring an IKE Access Profile</i> • Configuring an L2TP Access Profile on the LNS on page 246

interfaces (Static and Dynamic Subscribers)

```
Syntax  interfaces {
        interface-name {
            unit logical-unit-number {
                actual-transit-statistics;
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                    line-identity {
                        include {
                            accept-no-ids;
                            circuit-id;
                            remote-id;
                        }
                        dynamic-profile profile-name;
                    }
                }
            }
        }
        family family {
            access-concentrator name;
            address address;
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
                adf {
                    counter;
                    input-precedence precedence;
                    not-mandatory;
                    output-precedence precedence;
                    rule rule-value;
                }
                input filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
                output filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
            }
            max-sessions number;
            max-sessions-vsa-ignore;
            rpf-check {
                mode loose;
            }
            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;
            }
        }
    }
    service-name-table table-name
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
    unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name (
        precedence precedence;
        shared-name filter-shared-name;
    )
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    chap;
    pap;
}
proxy-arp;
service {
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
vlan-id;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
vlan-tagging;
}
interface-set interface-set-name {
    interface interface-name {
        unit logical unit number {
            advisory-options {
                downstream-rate rate;
                upstream-rate rate;
            }
        }
    }
}
pppoe-underlying-options {
    max-sessions number;
}
}
demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        family family {

```

```

access-concentrator name;
address address;
direct-connect;
duplicate-protection;
dynamic-profile profile-name;
demux-source {
    source-prefix;
}
filter {
    input filter-name (
        precedence precedence;
        shared-name filter-shared-name;
    )
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
mac-validate (loose | strict);
max-sessions number;
max-sessions-vsa-ignore;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name;
    output filter-name;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
pp0 {
    unit logical-unit-number {
        keepalives interval seconds;
        no-keepalives;
        pppoe-options {
            underlying-interface interface-name;
            server;
        }
        ppp-options {
            aaa-options aaa-options-name;
            authentication [ authentication-protocols ];
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            ignore-magic-number-mismatch;
            initiate-ncp (ip | ipv6 | dual-stack-passive)
            ipcp-suggest-dns-option;
            mru size;

```

```

mtu (size | use-lower-layer);
on-demand-ip-address;
pap;
peer-ip-address-optional;
}
family inet {
  unnumbered-address interface-name;
  address address;
  service {
    input {
      service-set service-set-name {
        service-filter filter-name;
      }
      post-service-filter filter-name;
    }
    output {
      service-set service-set-name {
        service-filter filter-name;
      }
    }
  }
}
filter {
  input filter-name {
    precedence precedence;
    shared-name filter-shared-name;
  }
  output filter-name {
    precedence precedence;
    shared-name filter-shared-name;
  }
}
}
}
}
}

```

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

Options *interface-name*—The interface variable (*\$junos-interface-ifd-name*). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.



NOTE: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *DHCP Subscriber Interface Overview*
- *Configuring Subscribers over Static Interfaces*
- *Demultiplexing Interface Overview*

ip-address-change-notify

Syntax	ip-address-change-notify { message <i>message</i> ; }
Hierarchy Level	[edit access profile <i>profile-name</i> radius options]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>For on-demand address allocation for dual-stack PPP subscribers, specify that the BNG includes the IPv4-Release-Control VSA (26–164) in the Access-Request that is sent during on-demand IP address allocation and in the Interim-Accounting messages that are sent to report an address change.</p> <p>The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.</p> <p>Optionally, configure a message that is included in the VSA when it is sent to the RADIUS server.</p>
Default	This functionality is disabled by default.
Options	<p>message—VSA message.</p> <p>Range: Up to 32 characters.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation</i> • <i>Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes</i> • <i>Enabling IPv4 Release Control VSA (26–164) in RADIUS Messages</i>

ip-reassembly

Syntax `ip-reassembly {
 profile profile-name
 rule rule-name{
 match-direction direction
 };
 }`

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the IP reassembly parameters to be applied to the L2TP server.



NOTE: Inline IP reassembly configuration does not require you to configure the **profile** statement. The **profile** configuration is used when IP reassembly is configured on services PICs.

Options **profile *profile-name***—Name of the IP reassembly profile.


The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring IP Inline Reassembly for L2TP on page 290](#)
- [IP Packet Fragment Reassembly for L2TP Overview on page 289](#)

ip-reassembly (L2TP)

Syntax	ip-reassembly { service-set <i>service-set-name</i> ; }
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Associate the reassembly service-set with the L2TP service.
<div>  NOTE: The service set must be defined at the [edit services] hierarchy level. </div>	
Options	service-set <i>service-set-name</i> —Identifies the service set to be associated with the L2TP service.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • IP Packet Fragment Reassembly for L2TP Overview on page 289 • Configuring IP Inline Reassembly for L2TP on page 290

ip-reassembly-rules (Service Set)

Syntax	<code>ip-reassembly-rules <i>rule-name</i>;</code>
Hierarchy Level	<code>[edit services service-setservice-set-name]</code>
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Specify one or more previously configured IP reassembly rules to associate with the service set.



NOTE: The IP reassembly rule must be defined at the `[edit services ip-reassembly rule]` hierarchy level.

Options	<i>rule-name</i> —Name of an IP reassembly rule.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IP Inline Reassembly for L2TP on page 290• IP Packet Fragment Reassembly for L2TP Overview on page 289

ipcp-suggest-dns-option

Syntax	<code>ipcp-suggest-dns-option;</code>
Hierarchy Level	<p>[edit access group-profile <i>group-profile-name</i> ppp ppp-options],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options],</p> <p>[edit interfaces pp0 unit <i>logical-unit-number</i> ppp-options],</p> <p>[edit interfaces <i>si-fpc/pic/port</i> unit <i>logical-unit-number</i> ppp-options]</p>
Release Information	Statement introduced in Junos OS Release 16.1.
Description	Configure the router to prompt Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for terminated PPPoE and LNS subscribers. You can configure this for dynamic or static PPPoE subscribers, dynamic or static LNS subscribers, and in an LNS group profile.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 136 • Configuring the Group Profile for Defining L2TP Attributes • Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245 • Configuring Dynamic Authentication for PPP Subscribers on page 125 • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

keepalive

Syntax	<code>keepalive seconds;</code>
Hierarchy Level	[edit access group-profile <i>profile-name</i> ppp], [edit access profile <i>profile-name</i> client <i>client-name</i> ppp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the keepalive interval for an L2TP tunnel.
Options	<p>seconds—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.</p> <p>For L2TP on MX Series routers, the minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.</p> <p>Range: 0 through 32,767 seconds</p> <p>Default: 30 seconds</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Group Profile for Defining L2TP Attributes</i>• <i>Configuring PPP Properties for a Client-Specific Profile</i>• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245

keepalives

Syntax	<code>keepalives <interval seconds> <down-count number> <up-count number>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.</p> <p>For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> • atm-ppp-llc—PPP over AAL5 LLC encapsulation. • atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Default	Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.
Options	<p>down-count <i>number</i>—The number of keepalive packets a destination must fail to receive before the network takes down a link.</p> <p>Range: 1 through 255</p> <p>Default: 3</p> <p>interval <i>seconds</i>—The time in seconds between successive keepalive requests.</p> <p>Range: 1 through 32767 seconds</p> <p>Default: 10 seconds</p> <p>up-count <i>number</i>—The number of keepalive packets a destination must receive to change a link's status from down to up.</p> <p>Range: 1 through 255</p> <p>Default: 1</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Keepalives • Configuring Frame Relay Keepalives • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

keepalives (Dynamic Profiles)

Syntax	<pre>keepalives { interval <i>seconds</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit <i>logical-unit-number</i>] [edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"] [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"] hierarchy level introduced in Junos OS Release 10.1. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"] hierarchy level introduced in Junos OS Release 12.2.
Description	Specify the keepalive interval in a PPP dynamic profile.
Default	Sending of keepalives is enabled by default.
Options	interval <i>seconds</i> —The time in seconds between successive keepalive requests. Range: 1 through 32767 seconds Default: 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Dynamic Profiles Overview• Configuring Dynamic Authentication for PPP Subscribers on page 125• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

l2tp

```
Syntax l2tp {
    access-line-information <connection-speed-update>;
    destination {
        address ip-address {
            access-line-information <connection-speed-update>;
            drain;
            routing-instance routing-instance-name {
                drain;
            }
        }
        lockout-timeout seconds;
        name destination-name {
            drain;
        }
    }
    destination-equal-load-balancing;
    destruct-timeout seconds;
    disable-calling-number-avp;
    disable-failover-protocol;
    drain;
    enable-ipv6-services-for-lac;
    enable-snmp-tunnel-statistics;
    failover-within-preference;
    ip-reassembly;
    maximum-sessions number;
    rx-connect-speed-when-equal;
    sessions-limit-group limit-group-name {
        maximum-sessions number;
    }
    traceoptions {
        debug-level level;
        file filename <files number> <match regular-expression> <size maximum-file-size>
            <world-readable | no-world-readable>;
        filter {
            protocol name;
            user user@domain;
            user-name username;
        }
        flag flag;
        interfaces interface-name {
            debug-level severity;
            flag flag;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    tunnel {
        assignment-id-format (assignment-id | client-server-id);
        idle-timeout seconds;
        maximum-sessions number;
        minimum-retransmission-timeout;
        name name {
```

```

    address ip-address {
        drain;
        routing-instance routing-instance-name {
            drain;
        }
    }
    drain;
}
nas-port-method;
retransmission-count-established count;
retransmission-count-not-established count;
rx-window-size packets;
tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port | reject |
    reject-ip-address | reject-udp-port);
}
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;
        gateway-name gateway-name;
    }
    maximum-send-window packets;
    maximum-sessions number;
    ppp-access-profile profile-name;
    receive-window packets;
    retransmit-interval seconds;
    service-device-pool pool-name;
    service-interface interface-name;
    service-profile profile-name (parameter) & profile-name;
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-value;
            services severity-level;
        }
    }
    tos-reflect;
    tunnel-switch-profile profile-name;
    tunnel-timeout seconds;
}
tunnel-switch-profile profile-name;
tx-connect-speed-method method;
weighted-load-balancing;
}

```

Hierarchy Level [edit services]

Release Information Statement 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 Configure L2TP services to establish PPP tunnels across a network.

The remaining statements are explained separately. See [CLI Explorer](#).



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Layer 2 Tunneling Protocol Overview](#)
- [L2TP for Subscriber Access Overview on page 147](#)

l2tp (Profile)

Syntax l2tp {
 interface-id *interface-id*;
 lcp-renegotiation;
 local-chap;
 maximum-sessions *number*;
 maximum-sessions-per-tunnel *number*;
 multilink {
 drop-timeout *milliseconds*;
 fragment-threshold *bytes*;
 }
 override-result-code session-out-of-resource;
 ppp-authentication (chap | pap);
 ppp-profile *profile-name*;
 sessions-limit-group;
 service-profile *profile-name*(*parameter*)&*profile-name*;
 shared-secret *shared-secret*;
 }

Hierarchy Level [edit access profile *profile-name* **client** *client-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the L2TP properties for a profile.

The remaining statements are explained separately.



NOTE: Only the **interface-id**, **lcp-renegotiation**, **maximum-sessions**, **maximum-sessions-per-tunnel**, **sessions-limit-group** and **shared-secret** statements are supported for L2TP LNS on MX Series routers.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring L2TP Properties for a Client-Specific Profile](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)

l2tp-access-profile

Syntax	l2tp-access-profile <i>profile-name</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the profile used to validate all L2TP connection requests to the local gateway address.
Options	<i>profile-name</i> —Identifier for the L2TP connection profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Access Profiles for L2TP Tunnel Groups</i>• Configuring an L2TP Access Profile on the LNS on page 246

l2tp-maximum-session (Service Interfaces)

Syntax	<code>l2tp-maximum-session <i>number</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>si-slot/pic/port</i>],</code> <code>[edit interfaces <i>asinumber</i>]</code>
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify the maximum number of L2TP sessions allowed on a physical service interface (si) or aggregated service interface (asi).</p> <p>New session requests on an interface are accepted only when the session count is less than the maximum session limit. If the limit has been reached, subsequent requests are dropped and the LNS responds with a CDN message (Result Code 2, Error Code 4). When a pool of interfaces is configured, interfaces at the maximum limit are ignored in favor of an interface in the pool that has a lower session count. For an asi interface, the configuration applies only to the asi interface. You cannot configure a session limit on the individual member interfaces of an asi bundle.</p> <p>Configuring the session limit to be less than the current number of sessions on the interface has no effect on existing sessions, but prevents any new sessions from being created until the number of session drops below the new limit.</p>
Options	<p><i>number</i>—Maximum number of L2TP sessions allowed for the interface. A value of 0 prevents the interface from being considered.</p> <p>Default: 64,000</p> <p>Range: 0 through 64,000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204• L2TP Session Limits and Load Balancing for Service Interfaces on page 259• Configuring an L2TP LAC on page 173• Configuring an L2TP LNS with Inline Service Interfaces on page 239• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274• L2TP for Subscriber Access Overview on page 147

layer2-liveness-detection (Receive)

Syntax	layer2-liveness-detection;
Hierarchy Level	[edit system services subscriber-management overrides interfaces family (inet inet6)]
Release Information	Statement introduced in Junos OS Release 17.4R1.
Description	<p>Enable a DHCP client host to determine the state of the DHCPv4 or DHCPv6 client session from the perspective of a router acting as a broadband network gateway (BNG). This statement causes the BNG to conduct a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients when it receives ARP or Neighbor Discovery (ND) packets.</p> <p>When the BNG receives either of these packets, it does the following:</p> <ol style="list-style-type: none"> 1. Checks whether Layer 2 liveness detection for subscriber management is enabled globally for the relevant address family, inet or inet6. 2. If liveness detection is not enabled, then the BNG responds as usual to the received packets without checking the state of the client session. If liveness detection is enabled for the family, then the BNG checks whether the client session is still in the bound state. 3. If the client session is bound, the BNG responds to the client with the appropriate ARP or ND packet. If the session is not bound, the BNG drops the received packet. It does not send an ARP or ND response packet to the host, enabling the host to determine that the BNG considers the session to be down. <p>This behavior can be referred to as the <i>receive</i> functionality for BNG Layer 2 liveness detection, as opposed to the <i>send</i> functionality configured with the layer2-liveness-detection (Send) statement for DHCP relay or DHCP local server.</p> <p>The usefulness of the receive functionality depends on the ability of the DHCP client host to reclaim resources from the stale client based on the absence of a response packet from the BNG for an unbound client session. If this capability requires a change in the client implementation, you may want to use the send functionality.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98 • DHCP Liveness Detection Overview on page 85 • <i>Configuring Junos OS Enhanced Subscriber Management</i>

layer2-liveness-detection (Send)


Syntax	<pre>layer2-liveness-detection { max-consecutive-retries <i>number</i>; transmit-interval <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method], [edit forwarding-options dhcp-relay dual-stack-group <i>dual-stack-group-name</i> liveness-detection method], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method], [edit forwarding-options dhcp-relay liveness-detection method],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method], [edit system services dhcp-local-server dhcpv6 liveness-detection method], [edit system services dhcp-local-server <i>dual-stack-group</i> <i>dual-stack-group-name</i> liveness-detection method], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method], [edit system services dhcp-local-server liveness-detection method],</p>
Release Information	Statement introduced in Junos OS Release 17.4R1.
Description	<p>Configure a router acting as a broadband network gateway (BNG) to conduct a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients to determine the validity and state of the DHCP client session, and to clean up inactive sessions.</p> <p>The BNG sends ARP or ND request packets to the each DHCP client at a configurable interval, then waits for a response. If it receives a response from a client before the interval times out, it sends another request to the client when the timer expires.</p> <p>If the BNG does not receive a response before the interval times out, it sets the timer to 30 seconds and sends another request. This is the first retry attempt.</p> <p>If it receives a response from a client before the 30-second interval times out, it sends another request to the client when the timer expires. If the 30-second timer expires before a response is received, the BNG sets the timer to 10 seconds and sends another request. This is the second retry attempt. If the BNG does not receive a response within this interval it resets the timer to 10 seconds and sends another request. The BNG continues to send requests at 10-second intervals until it either receives a response from the client before the interval times out or exhausts the number of retry attempts.</p> <p>The first retry attempt uses a 30-second interval. Subsequent retries occur at 10-second intervals. The number of possible 10-second retries is therefore the total number of retries minus 1. For example, if you configure 5 retries, there is one 30-second retry and up to four 10-second retries.</p> <p>If the BNG attempts all the retries and never receives a response from a client within the interval, the client session is declared to be down.</p>



NOTE: The only option to the **failure-action** statement supported by Layer 2 liveness detection is **clear-binding**.

Options	<p>max-consecutive-retries <i>number</i>—Maximum number of consecutive times that the router sends an ARP request packet in the absence of an ARP response packet. Range: 3 through 6 retries Default: 3 retries</p> <p>transmit-interval <i>seconds</i>—Initial interval that the router waits for an ARP response after sending an ARP request packet to the client or waits for an ND response packet after sending an NG request packet to the client. Range: 300 through 1800 seconds Default: 300 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98 • DHCP Liveness Detection Overview on page 85

lcp-renegotiation

Syntax	lcp-renegotiation;
Hierarchy Level	[edit access group-profile <i>profile-name</i> l2tp], [edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client. When LCP renegotiation is disabled, LNS uses the pre-negotiated LCP parameters between the L2TP access concentrator (LAC) and PPP client to set up the session. When LCP renegotiation is enabled, authentication is also renegotiated.
<div> NOTE: This statement is not supported at the [edit access group-profile l2tp] hierarchy level for L2TP LNS on MX Series routers.</div>	
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring L2TP for a Group Profile</i>• <i>Configuring L2TP Properties for a Client-Specific Profile</i>• Configuring an L2TP Access Profile on the LNS on page 246

liveness-detection

Syntax

```
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
  }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay group group-name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server group group-name]
```

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure bidirectional failure detection timers and authentication criteria for static routes.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.


Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

- Related Documentation**
- [DHCP Liveness Detection Overview on page 85](#)
 - [Configuring Detection of DHCP Local Server Client Connectivity with BFD on page 92](#)
 - [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD on page 87](#)
 - [Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94](#)
 - [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)
 - [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)

local-gateway (L2TP LNS)

Syntax	<pre>local-gateway { address address; gateway-name gateway-name; }</pre>
Hierarchy Level	[edit services l2tp tunnel-group name]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Options	<p>address—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.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Local Gateway Address and PIC.• Configuring L2TP Tunnel Groups• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274

lockout-timeout (L2TP Destination Lockout)

Syntax	<code>lockout-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit services l2tp destination]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Set the duration of the timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created. Destinations are locked out when L2TP cannot connect to the destination during the tunnel selection process. This statement does not affect destinations that are currently locked out.
<div>  NOTE: The <i>ip-address</i> option for the <i>destination</i> statement does not apply to the <i>lockout-timeout</i> statement. </div>	
Options	<p><i>seconds</i>—Length of the period during which the destination is locked out.</p> <p>Range: 60 through 3600 seconds</p> <p>Default: 300 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the L2TP Destination Lockout Timeout on page 165 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

logical-system (Tunnel Profile)

Syntax	<code>logical-system <i>logical-system-name</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify a logical system for a tunnel. When you specify a logical system, you must also specify a routing instance.
Options	<i>logical-system-name</i> — Name of the logical system. Default: Logical system <i>default</i>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

mac

Syntax	<code>mac <i>mac-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the MAC address of the interface. Use this statement at the [edit interfaces ... ps0] hierarchy level to configure the MAC address for a pseudowire logical device that is used for subscriber interfaces over point-to-point MPLS pseudowires.
Options	<i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0000.5e00.5355 or 00:00:5e:00:53:55.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the MAC Address on the Management Ethernet Interface• Configuring a Pseudowire Subscriber Logical Interface Device on page 315

mac-address (Dynamic Access-Internal Routes)

Syntax	<code>mac-address address;</code>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access-internal route <i>subscriber-ip-address</i> qualified-next-hop <i>underlying-interface</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access-internal route <i>subscriber-ip-address</i> qualified-next-hop <i>underlying-interface</i>],</p> <p>[edit dynamic-profiles routing-options access-internal route <i>subscriber-ip-address</i> qualified-next-hop <i>underlying-interface</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options route <i>subscriber-ip-address</i> qualified-next-hop <i>underlying-interface</i>] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> route <i>subscriber-ip-address</i> qualified-next-hop <i>underlying-interface</i>] hierarchy levels introduced in Junos OS Release 10.1.</p>
Description	Dynamically configure the MAC address variable for an access-internal route for unnumbered interfaces such as DHCP subscriber interfaces.
Options	<i>address</i> —Either the specific MAC address you want to assign to the access-internal route or the MAC address variable (\$junos-subscriber-mac-address). The MAC address variable is dynamically replaced with the value supplied by DHCP when a subscriber logs in.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32

match-direction (IP Reassembly Rule)

Syntax	<code>match-direction <i>direction</i></code>
Hierarchy Level	[edit services ip-reassemblyrule <i>rule-name</i>]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Configure the direction in which the IP reassembly rule matching is applied. The match direction is used with respect to the traffic flow through the inline services interface. You must configure a match direction for an IP reassembly rule.
Options	<i>direction</i> —Match direction. For inline IP reassembly, input is the only match direction supported.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IP Inline Reassembly for L2TP on page 290• IP Packet Fragment Reassembly for L2TP Overview on page 289

maximum-sessions (L2TP)

Syntax	<code>maximum-sessions <i>number</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> <i>client</i> <i>client-name</i>],</code> <code>[edit services <i>l2tp</i>],</code> <code>[edit services <i>l2tp</i> <i>sessions-limit-group</i>],</code> <code>[edit services <i>l2tp</i> <i>tunnel</i>],</code> <code>[edit services <i>l2tp</i> <i>tunnel-group</i> <i>group-name</i>],</code>
Release Information	Statement introduced in Junos OS Release 16.1.
Description	Specify the maximum number of L2TP sessions for the chassis, all tunnels, a tunnel group, a session limit group, or a client.
Options	<p><i>number</i>—Number of sessions allowed.</p> <p>Range: (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit</p> <p>Range: (Tunnel) 1 through 65,536</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239 • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274 • L2TP for Subscriber Access Overview on page 147

maximum-sessions-per-tunnel

Syntax	<code>maximum-sessions-per-tunnel <i>number</i>;</code>
Hierarchy Level	<code>[edit access group-profile l2tp],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the maximum sessions for a Layer 2 tunnel.



NOTE: This statement is not supported at the `[edit access group-profile l2tp]` hierarchy level for L2TP LNS on MX Series routers.

Options	<i>number</i> —Maximum number of sessions for a Layer 2 tunnel.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP for a Group Profile• Configuring L2TP Properties for a Client-Specific Profile• Configuring an L2TP Access Profile on the LNS on page 246

max-sessions (Tunnel Profile)

Syntax	<code>max-sessions <i>number</i>;</code>
Hierarchy Level	<code>[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the maximum number of sessions allowed in the tunnel.
Options	<p><i>number</i>—Maximum number of sessions allowed in the tunnel. A value of 0 means that the maximum configurable number of sessions is allowed.</p> <p>Range: 0 through 60,000</p> <p>Default: 0</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Tunnel Profile for Subscriber Access on page 207

medium (Tunnel Profile)

Syntax	<code>medium <i>type</i>;</code>
Hierarchy Level	<code>[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the medium type for the tunnel.
Default	ipv4
Options	<p><i>type</i>—Medium type for the tunnel. The only value currently available is ipv4.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Tunnel Profile for Subscriber Access on page 207

method

```
Syntax  method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
```

Hierarchy Level [edit forwarding-options dhcp-relay dhcpv6 group *group-name* [liveness-detection](#)],
 [edit forwarding-options dhcp-relay dhcpv6 [liveness-detection](#)],
 [edit forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name* [liveness-detection](#)],
 [edit forwarding-options dhcp-relay group *group-name* [liveness-detection](#)],
 [edit forwarding-options dhcp-relay [liveness-detection](#)],
 [edit system services dhcp-local-server dhcpv6 group *group-name* [liveness-detection](#)],
 [edit system services dhcp-local-server dhcpv6 [liveness-detection](#)],
 [edit system services dhcp-local-server dual-stack-group *dual-stack-group-name* [liveness-detection](#)],
 [edit system services dhcp-local-server group *group-name* [liveness-detection](#)],
 [edit system services dhcp-local-server [liveness-detection](#)]

Release Information Statement introduced in Junos OS Release 12.1.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the liveness detection method.



NOTE: The [bfd](#) stanza is not available at the [edit forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name* [liveness-detection](#) [method](#)] or [edit system services dhcp-local-server dual-stack-group *dual-stack-group-name* [liveness-detection](#)] hierarchy levels.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• DHCP Liveness Detection Overview on page 85• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89• DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98
------------------------------	---

metric (Dynamic Access-Internal Routes)

Syntax	<code>metric route-cost;</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>],</code> <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>],</code> <code>[edit dynamic-profiles <i>profile-name</i> routing-options access route <i>prefix</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5. Support at the <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> and <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> hierarchy levels introduced in Junos OS Release 10.1.
Description	Dynamically configure the cost for an access route.
Options	<code>route-cost</code> —Either the specific cost you want to assign to the access route or either of the following cost variables: <ul style="list-style-type: none">• <code>\$junos-framed-route-cost</code>—Cost of an IPv4 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-Route attribute [22].• <code>\$junos-framed-route-ipv6-cost</code>—Cost of an IPv6 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-IPv6-Route attribute [99].
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Dynamic Access Routes for Subscriber Management on page 33

minimum-interval

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method <code>bfd</code>], [edit system services dhcp-local-server liveness-detection method bfd <code>transmit-interval</code>], [edit system services dhcp-local-server dhcpv6 liveness-detection method <code>bfd</code>], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <code>transmit-interval</code>], [edit forwarding-options dhcp-relay liveness-detection method <code>bfd</code>], [edit forwarding-options dhcp-relay liveness-detection method bfd <code>transmit-interval</code>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <code>bfd</code>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <code>transmit-interval</code>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <code>bfd</code>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <code>transmit-interval</code>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <code>bfd</code>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <code>transmit-interval</code>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <code>bfd</code>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <code>transmit-interval</code>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <code>bfd</code>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <code>transmit-interval</code>]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <code>transmit-interval</code> <code>minimal-interval</code> and <code>minimum-receive-interval</code> statements.</p>
Options	<p><i>milliseconds</i> — Specify the minimum interval value for BFD liveliness detection.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94

- [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)

minimum-receive-interval

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit system services dhcp-local-server liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options</code> <code> dhcp-relay dhcpv6 liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code> bfd],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code> bfd]</code>
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.
Options	<i>milliseconds</i> — Specify the minimum receive interval value. Range: 1 through 255,000
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

minimum-retransmission-timeout (L2TP Tunnel)

Syntax	minimum-retransmission-timeout <i>seconds</i> ;
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Configure the minimum (initial) interval that the LAC or the LNS waits for a response after transmitting an L2TP control message to a peer. If no response has been received by the time the period expires, the message is retransmitted. The timeout period is doubled for each retransmission until the maximum of 16 seconds is reached.
Options	<i>seconds</i> —Minimum interval before initial retransmission. Range: 1, 2, 4, 8, or 16 seconds Default: 1
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Retransmission Attributes for L2TP Control Messages on page 171 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

mtu

Syntax `mtu bytes;`

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number family family],
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols l2vpn interface interface-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit logical-systems name protocols ospf area name interface ],
[edit logical-systems name routing-instances name protocols ospf area name interface],
[edit protocols ospf area name interface ],
[edit routing-instances name protocols ospf area name interface]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Support at the `[set interfaces interface-name unit logical-unit-number family ccc]` hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.

Statement introduced in Junos OS 17.3R1 Release for MX Series Routers.

Description Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named `irb` or `vlan`, respectively).



.....

CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

.....



.....

NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the `flexible-vlan-tagging` statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
 - In case of Layer 2 IFL configured with the `vlan-tagging` statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.
-



NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the `mtu` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] or [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
 - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
 - If you are configuring the protocol MTU for both inet and inet6 families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for inet and inet6 families that are configured on the same logical interface.
- Starting in Release 14.2, MTU for IRB interfaces is calculated by removing the Ethernet header overhead (6(DMAC)+6(SMAC)+2(EtherType)), and the MTU is a minimum of the two values:
 - Configured MTU
 - Associated bridge domain's physical or logical interface MTU
 - For Layer 2 logical interfaces configured with flexible-vlan-tagging, IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
 - For Layer 2 logical interfaces configured with vlan-tagging, IRB MTU is calculated by including single VLAN 4 bytes overhead.



NOTE: Changing the Layer 2 logical interface option from `vlan-tagging` to `flexible-vlan-tagging` or vice versa adjusts the logical interface MTU by 4 bytes with the existing MTU size. As a result, the Layer 2 logical interface is deleted and re-added, and the IRB MTU is re-computed appropriately.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers), 256 through 9500 bytes (Junos OS 16.1R1 for MX Series routers)



NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Media MTU*
- *Configuring the MTU for Layer 2 Interfaces*
- *Setting the Protocol MTU*

multiplier

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.
Options	number —Maximum allowable number of hello packets missed by the neighbor. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

name (L2TP Destination)

Syntax	<code>name <i>destination-name</i> { drain; }</code>
Hierarchy Level	[edit services l2tp destination]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the name of the L2TP destination for the tunnel.
Options	<i>destination-name</i> —Locally assigned name of the tunnel destination. The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Drain on page 166

name (L2TP Tunnel Destination)

Syntax `name name {
 address ip-address {
 drain;
 routing-instance routing-instance-name {
 drain;
 }
 }
 drain;
 }`

Hierarchy Level [edit services l2tp [tunnel](#)]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the local name and other attributes of the L2TP tunnel.

Options *name* —Locally assigned name of the tunnel; in the format *destination-name/tunnel-name* or *tunnel-name*.



NOTE: When only the tunnel name is provided, then you must identify the destination for the tunnel by including the `address ip-address` statement at the [edit services l2tp tunnel name *name*] hierarchy level.

The remaining statements are explained separately. See [CLI Explorer](#).


Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring L2TP Drain on page 166](#)

no-adaptation

Syntax	no-adaptation;
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

nas-port-method (L2TP LAC)

Syntax	nas-port-method cisco-avp;
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Globally configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.
<div> NOTE: This global configuration can be overridden by the configuration in a tunnel profile or by the RADIUS configuration.</div>	
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Globally Configuring the LAC to Interoperate with Cisco LNS Devices on page 178• Configuring a Tunnel Profile for Subscriber Access on page 207• LAC Interoperation with Third-Party LNS Devices on page 177


nas-port-method (Tunnel Profile)

Syntax	nas-port-method cisco-avp;
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207


next-hop (Dynamic Access-Internal Routes)

Syntax	<code>next-hop <i>next-hop</i>;</code>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access <i>route prefix</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access <i>route prefix</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options <i>access route prefix</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access <i>route prefix</i>] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access <i>route prefix</i>] hierarchy levels introduced in Junos OS Release 10.1.</p>
Description	<p>Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.</p> <p>The next-hop gateway can be specified explicitly in the framed route, as either the subscriber's fixed address (common for business subscribers) or 0.0.0.0. Alternatively, the absence of the gateway address implies address 0.0.0.0. The address 0.0.0.0, whether implicit or explicitly configured, resolves to the subscriber's assigned address (host route).</p> <p>If the RADIUS Framed-Route attribute [22] or Framed-IPv6-Route attribute [99] does not specify the next-hop gateway—as is common—the variable representing the next-hop automatically resolves to the subscriber's IP address.</p>
Options	<p><i>next-hop</i>—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.</p> <ul style="list-style-type: none"> For IPv4 access routes, use the variable, \$junos-framed-route-nexthop. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22]. For IPv6 access routes, use the variable, \$junos-framed-route-ipv6-nexthop. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Dynamic Access Routes for Subscriber Management on page 33

next-hop-service

Syntax	<pre> next-hop-service { inside-service-interface <i>interface-name.unit-number</i>; outside-service-interface <i>interface-name.unit-number</i>; outside-service-interface-type <i>interface-type</i>; service-interface-pool <i>name</i>; } </pre>
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>service-interface-pool option added in Junos OS Release 9.3.</p>
Description	Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.
Options	<p>inside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied inside the network.</p> <p>outside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied outside the network.</p> <p>outside-service-interface-type <i>interface-type</i>—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.</p> <p>service-interface-pool <i>name</i>—Name of the pool of logical interfaces configured at the [edit services service-interface-pools pool <i>pool-name</i>] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.</p>
<div>  <p>NOTE: service-interface-pool is not applicable for IP reassembly configuration on L2TP.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Service Sets to be Applied to Services Interfaces

no-allow-snooped-clients

Syntax	no-allow-snooped-clients;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Explicitly disable DHCP snooping support on DHCP relay agent.</p> <p>Use the statement at the [edit ... dhcpv6] hierarchy levels to explicitly disable snooping support for DHCPv6 relay agent.</p>
<div>  <p>NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is <i>enabled</i> by default. In Release 10.1 and later, DHCP snooping is <i>disabled</i> by default.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Extended DHCP Relay Agent Overview</i> • <i>Overriding the Default DHCP Relay Configuration Settings</i> • DHCP Snooping Support on page 41 • Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48

no-gratuitous-arp-request

Syntax	no-gratuitous-arp-request;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.
Default	Gratuitous ARP responses are enabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gratuitous ARP</i>

no-snoop (DHCP Local Server and Relay Agent)

Syntax	no-snoop;
Hierarchy Level	[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ...], [edit routing-instances <i>routing-instance-name</i> ...], [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6]
Release Information	Statement introduced in Junos OS Release 15.1R2.
Description	<p>Disable DHCP snooping filters.</p> <p>DHCP snooping provides DHCP security by identifying incoming DHCP packets. In the default DHCP snooping configuration, all traffic is snooped. You can optionally use the forward-snooped-clients statement to evaluate the snooped traffic and to determine if the traffic is forwarded or dropped, based on whether or not the interface is configured as part of a group.</p> <p>In both the default configuration and in configurations using the forward-snooped-clients statement, all DHCP traffic is forwarded from the hardware control plane to the routing plane of the routing instance to ensure that all DHCP packets are intercepted. In certain topologies, such as a Metropolitan Routing Ring topology, forwarding all DHCP traffic to the control plane can result in excessive traffic. The no-snoop configuration statement disables the snooping filter for DHCP traffic that can be directly forwarded on the hardware control plane, such as Layer 3 unicast packets with a valid route, causing those DHCP packets to bypass the slower routing plane.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling DHCP Snooping Filters on page 51 • Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 43 • Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 48

no-vlan-id-validate

Syntax	no-vlan-id-validate;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Uniquely identify a Layer 2 circuit for either a standard pseudowire or a redundant pseudowire.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for Layer 2 Circuits• Pseudowire Subscriber Logical Interfaces Overview on page 307• Configuring a Pseudowire Subscriber Logical Interface on page 313• Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces on page 318

on-demand-ip-address

Syntax	on-demand-ip-address;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options], [edit interfaces pp0 unit <i>unit-number</i> ppp-options], [edit protocols ppp-service]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>For IPv4 and IPv6 dual-stack PPP subscribers, enables on-demand allocation and de-allocation of an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address.</p> <p>Configuration changes take effect as follows:</p> <ul style="list-style-type: none"> • When you change this setting for a dynamic PPP interface (at the [edit dynamic-profiles] hierarchy level), the change takes effect only for new subscriber logins. • When you change this setting for a static PPP interface (at the [edit interfaces pp0] hierarchy level), the subscribers on the interface are logged out. • When you change this setting globally (at the [edit protocols ppp-service] hierarchy level), the change takes effect only for new subscriber logins. <p>If you enable on-demand allocation at both the interface and global levels, the global configuration takes precedence and changes take effect for new subscriber logins.</p>
Default	This functionality is disabled by default.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation</i> • <i>Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers</i> • <i>Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers</i> • <i>Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers</i>

overrides (DHCP Relay Agent)

Syntax	<pre> overrides { allow-no-end-option; allow-snooped-clients; always-write-giaddr; always-write-option-82; asymmetric-lease-time <i>seconds</i>; asymmetric-prefix-lease-time <i>seconds</i>; client-discover-match <option60-and-option82 incoming-interface>; client-negotiation-match incoming-interface; delay-authentication; delete-binding-on-renegotiation; disable-relay; dual-stack <i>dual-stack-group-name</i>; interface-client-limit <i>number</i>; layer2-unicast-replies; no-allow-snooped-clients; no-bind-on-request; proxy-mode; relay-source replace-ip-source-with; send-release-on-delete; trust-option-82; } </pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Override the default configuration settings for the extended DHCP relay agent. Specifying the overrides statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p> <p>The following statements are supported at both the [edit ... dhcp-relay] and [edit ... dhcpv6] hierarchy levels.</p> <ul style="list-style-type: none"> allow-snooped-clients

- `asymmetric-lease-time`
- `delete-binding-on-renegotiation`
- `dual-stack`
- `interface-client-limit`
- `no-allow-snooped-clients`
- `no-bind-on-request`
- `relay-source`
- `send-release-on-delete`

The following statements are supported at the `[edit ... dhcpv6]` hierarchy levels only.

- `asymmetric-prefix-lease-time`

All other statements are supported at the `[edit ... dhcp-relay]` hierarchy levels only.

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	• <i>Extended DHCP Relay Agent Overview</i>
	• <i>Overriding the Default DHCP Relay Configuration Settings</i>
	• <i>Deleting DHCP Local Server and DHCP Relay Override Settings</i>

overrides (Enhanced Subscriber Management)

Syntax

```
overrides {
  interfaces {
    family (inet | inet6) {
      layer2-liveness-detection;
    }
  }
  no-unsolicited-ra;
  shmlog {
    disable;
    file filename <files maximum-no-files> <size maximum-file-size>;
    filtering enable;
    log-name {
      all;
      logname {
        <brief | detail | extensive | none | terse>;
        <file-logging | no-file-logging>;
      }
    }
    log-type (debug | info | notice);
  }
}
```

Hierarchy Level [edit system services [subscriber-management](#)]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Override the default configuration settings for the Junos OS enhanced subscriber management software for subscriber management.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.


Related Documentation

- [Configuring Junos OS Enhanced Subscriber Management](#)
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)

override-result-code (L2TP Profile)

Syntax	<pre>override-result-code { session-out-of-resource; }</pre>
Hierarchy Level	[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure the LNS to override result codes in Call-Disconnect-Notify (CDN) messages.
Options	session-out-of-resource —Override result codes 4 and 5 with result code 2. These result codes indicate that the number of L2TP sessions have reached the configured maximum value and the LNS can support no more sessions. When the LAC receives the code, it fails over to another LNS to establish subsequent sessions. Some third-party LACs respond only to result code 2.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an L2TP Access Profile on the LNS on page 246

pap

Syntax	<pre>pap { access-profile <i>name</i>; default-pap-password <i>password</i>; local-name <i>name</i>; local-password <i>password</i>; passive; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> ppp-options],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]</p>
Release Information	Statement introduced in Junos OS Release 8.3.
Description	<p>Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.</p> <p>After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.</p>
	<p> BEST PRACTICE: On inline service (si) interfaces for L2TP, only the <code>pap</code> statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.</p>
	<p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the PPP Challenge Handshake Authentication Protocol</i> • <i>Configuring the PPP Password Authentication Protocol On a Logical Interface</i> • <i>Tracing Operations of the pppd Process</i> • <i>traceoptions (PPP Process)</i> • <i>Example: Configuring PAP for an L2TP Profile</i> • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

pap (Dynamic PPP)

Syntax	<code>pap;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options] hierarchy level introduced in Junos OS Release 12.2.
Description	Specify PAP authentication in a PPP dynamic profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Dynamic Profiles Overview • Configuring Dynamic Authentication for PPP Subscribers on page 125 • Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121 • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

pap (L2TP)

Syntax	<code>pap;</code>
Hierarchy Level	[edit access group-profile <i>profile-name</i> ppp ppp-options]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	(MX Series routers only) Specify PAP authentication for PPP subscribers in an L2TP LNS user group profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245

parse-direction (Access Profile)

Syntax	<code>parse-direction (left-to-right right-to-left);</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> session-options strip-user-name]</code>
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify the direction in which a subscriber login string is parsed to identify the first delimiter that matches one configured with the delimiter statement. When subscriber username stripping is configured in a subscriber access profile, the characters to the right of the identified delimiter are stripped and discarded along with the delimiter. The remaining characters become the new, modified username.</p>
Default	<code>left-to-right</code>
Options	<p>left-to-right—Parse the subscriber login string from left to right up to the delimiter.</p> <p>For example, when the direction is left-to-right, the characters <code>/@\$%#</code> are configured as the delimiters, and the login string is <code>drgt21@example.com\$84</code>, the <code>@</code> is reached before the <code>\$</code>, so the username is modified to <code>drgt21</code>.</p> <p>right-to-left—Parse the subscriber login string from right to left up to the delimiter.</p> <p>For example, when the direction is right-to-left, the characters <code>/@\$%#</code> are configured as the delimiters, and the login string is <code>drgt21@example.com\$84</code>, the <code>\$</code> is reached before the <code>@</code>, so the username is modified to <code>drgt21@example.com</code>.</p>
Required Privilege Level	<p><code>access</code>—To view this statement in the configuration.</p> <p><code>access-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Session Options for Subscriber Access• Configuring Username Modification for Subscriber Sessions on page 179• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

pic (M Series and T Series Routers)

```
Syntax  pic pic-number {
        cel {
            el port-number {
                channel-group group-number timeslots slot-number;
            }
        }
        ct3 {
            port port-number {
                tl link-number {
                    channel-group group-number timeslots slot-number;
                }
            }
        }
        framing (sdh | sonet);
        idle-cell format {
            itu-t;
            payload-pattern payload-pattern-byte;
        }
        inline-services {
            bandwidth (1g | 10g);
        }
        max-queues-per-interface (8 | 4);
        no-concatenate;
    }
```

Hierarchy Level [edit chassis fpc *slot-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure properties for an individual PIC.

Options *pic-number*—Slot number in which the PIC is installed.
Range: 0 through 3

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 the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs*
- *Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode*
- *Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots*
- *Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs*

pool (L2TP Service Interfaces)

Syntax	<code>pool <i>pool-name</i> { interface <i>service-interface-name</i>; }</code>
Hierarchy Level	[edit services service-device-pools]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Define a pool of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.
Options	<p><i>pool-name</i>—Name of the service interface pool.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 284

pp0 (Dynamic PPoE)

```

Syntax  pp0 {
        unit logical-unit-number {
            keepalives interval seconds;
            no-keepalives;
            pppoe-options {
                underlying-interface interface-name;
                server;
            }
            ppp-options {
                aaa-options aaa-options-name;
                authentication [ authentication-protocols ];
                chap {
                    challenge-length minimum minimum-length maximum maximum-length;
                }
                ignore-magic-number-mismatch;
                initiate-ncp (ip | ipv6 | dual-stack-passive)
                ipcp-suggest-dns-option;
                mru size;
                mtu (size | use-lower-layer);
                on-demand-ip-address;
                pap;
                peer-ip-address-optional;
            }
            family inet {
                unnumbered-address interface-name;
                address address;
                service {
                    input {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                filter {
                    input filter-name {
                        precedence precedence;
                    }
                    output filter-name {
                        precedence precedence;
                    }
                }
            }
        }
    }

```

Hierarchy Level [edit dynamic-profiles *profile-name* [interfaces](#)]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the dynamic PPPoE logical interface in a dynamic profile. When the router creates a dynamic PPPoE logical interface on an underlying Ethernet interface configured with PPPoE (**ppp-over-ether**) encapsulation, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring a PPPoE Dynamic Profile*
- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- For information about creating static PPPoE interfaces, see *Configuring PPPoE*

ppp (Group Profile)

Syntax

```
ppp {
  cell-overhead;
  encapsulation-overhead bytes;
  framed-pool framed-pool;
  idle-timeout seconds;
  interface-id interface-id;
  keepalive seconds;
  ppp-options {
    aaa-options aaa-options-name;
    chap;
    ignore-magic-number-mismatch;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    mru;
    mtu;
    pap;
    peer-ip-address-optional;
  }
  primary-dns primary-dns;
  primary-wins primary-wins;
  secondary-dns secondary-dns;
  secondary-wins secondary-wins;
}
```

Hierarchy Level [edit access [group-profile](#) *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure PPP properties for a group profile.
The remaining statements are explained separately.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Group Profile for Defining L2TP Attributes](#)
- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)

ppp-options

Syntax	<pre> ppp-options { authentication [authentication-protocols]; mru size; mtu (size use-lower-layer); chap { access-profile name; challenge-length minimum minimum-length maximum maximum-length; default-chap-secret name; local-name name; passive; } compression { acfc; pfc; } dynamic-profile profile-name; initiate-ncp (ip ipv6 dual-stack-passive) ipcp-suggest-dns-option; lcp-max-conf-req number lcp-restart-timer milliseconds; loopback-clear-timer seconds; ncp-max-conf-req number ncp-restart-timer milliseconds; on-demand-ip-address pap { access-profile name; default-pap-password password; local-name name; local-password password; passive; } } </pre>
Hierarchy Level	<pre> [edit interfaces interface-name], [edit interfaces interface-name unit logical-unit-number], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number] </pre>

Release Information Statement introduced before Junos OS Release 7.4.

Description On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.



BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)

ppp-options (Dynamic PPP)

Syntax	<pre>ppp-options { aaa-options <i>aaa-options-name</i>; authentication [<i>authentication-protocols</i>]; chap { challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>; } ignore-magic-number-mismatch; initiate-ncp (ip ipv6 dual-stack-passive) ipcp-suggest-dns-option; mru <i>size</i>; mtu (<i>size</i> use-lower-layer); on-demand-ip-address; pap; peer-ip-address-optional; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"] hierarchy level introduced in Junos OS Release 12.2.
Description	Configure PPP-specific interface properties in a dynamic profile. The remaining statements are explained separately. See CLI Explorer .



NOTE:

PPP options can also be configured in a group profile with the [ppp-options \(L2TP\)](#) statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring a PPPoE Dynamic Profile*
- [Configuring Dynamic Authentication for PPP Subscribers on page 125](#)
- [Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 121](#)
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242](#)

ppp-options (L2TP)

Syntax `ppp-options {
 aaa-options aaa-options-name;
 chap;
 ignore-magic-number-mismatch;
 initiate-ncp (ip | ipv6 | dual-stack-passive)
 ipcp-suggest-dns-option;
 mru;
 mtu;
 pap;
 peer-ip-address-optional;
 }`

Hierarchy Level [edit access group-profile *profile-name* **ppp**]

Release Information Statement introduced in Junos OS Release 11.4.
 mtu statement introduced in Junos OS Release 14.2

Description Configure PPP-specific properties in a group profile that applies to tunneled PPP subscribers at the LNS.

The remaining statements are explained separately. See [CLI Explorer](#).



NOTE:

PPP options can also be configured for an inline service interface within a dynamic profile with the **ppp-options (Dynamic PPP)** statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both the dynamic profile and the group profile, the group profile options are applied to the subscriber only when the dynamic profile PPP options do not include any of the following attributes: **aaa-options**, **chap**, **ipcp-suggest-dns-option**, **mru**, **mtu**, **pap**, and **peer-ip-address-optional**. When any of these attributes is present, the dynamic profile is applied to the subscriber.

When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Group Profile for Defining L2TP Attributes](#)
- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245](#)

preference (Subscriber Management)

Syntax `preference route-distance`

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options access [route prefix](#)],
 [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options rib *routing-table-name* access [route prefix](#)],
 [edit dynamic-profiles *profile-name* routing-options [access route prefix](#)]

Release Information Statement introduced in Junos OS Release 9.5.
 Support at [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options access *route prefix*] and [edit dynamic-profiles *profile-name* routing-instances \$junos-routing-instance routing-options rib *routing-table-name* access *route prefix*] hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the distance for an access route.

Options *route-distance*—Either the specific distance you want to assign to the access route or either of the following distance variables:

- **\$junos-framed-route-distance**—Distance of an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-Route attribute [22].
- **\$junos-framed-route-ipv6-distance**—Distance of an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)

preference (Tunnel Profile)

Syntax	<code>preference <i>number</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>Specify the preference for a tunnel. You can specify up to 8 levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.</p> <p>This value can be overridden by RADIUS attribute Tunnel-Preference [83].</p>
Options	<p><i>number</i>—Number that indicates the order in which the router attempts to connect to the destination. Zero is the highest level of preference.</p> <p>Range: 0 through 2000</p> <p>Default: 2000</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

primary-interface (Aggregated Inline Services)

Syntax	<code>primary-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit interfaces asix aggregated-inline-services-options]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify the primary (active) inline services member link in the asi bundle. You must also configure a secondary (backup) member link on a different MPC with the secondary-interface statement. The secondary member provides 1:1 redundancy for subscriber service sessions on the primary member link. The bandwidth configured at the <code>[edit chassis fpc slot pic number inline-services bandwidth]</code> hierarchy level must be the same for both member links.</p> <p>Redundancy is not achievable if you configure the primary and secondary interface on the same MPC, because both member interfaces go down if the card goes down. Consequently, if you configure both interfaces on the same MPC, the subsequent configuration commit fails.</p>
Options	<i>interface-name</i> —Name of an inline services physical interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

profile (Access)

```
Syntax  profile profile-name {
        accounting {
            address-change-immediate-update
            accounting-stop-on-access-deny;
            accounting-stop-on-failure;
            ancp-speed-change-immediate-update;
            coa-immediate-update;
            coa-no-override service-class-attribute;
            duplication;
            duplication-filter;
            duplication-vrf {
                access-profile-name profile-name;
                vrf-name vrf-name;
            }
            immediate-update;
            order [ accounting-method ];
            send-acct-status-on-config-change;
            statistics (time | volume-time);
            update-interval minutes;
            wait-for-acct-on-ack;
        }
        accounting-order (radius | [accounting-order-data-list]);
        authentication-order [ authentication-methods ];
        client client-name {
            chap-secret chap-secret;
            group-profile profile-name;
            ike {
                allowed-proxy-pair {
                    remote remote-proxy-address local local-proxy-address;
                }
                pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
                ike-policy policy-name;
                interface-id string-value;
            }
        }
        l2tp {
            aaa-access-profile profile-name;
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions number;
            maximum-sessions-per-tunnel number;
            multilink {
                drop-timeout milliseconds;
                fragment-threshold bytes;
            }
            override-result-code session-out-of-resource;
            ppp-authentication (chap | pap);
            ppp-profile profile-name;
            service-profile profile-name(parameter)&profile-name;
            sessions-limit-group limit-group-name;
            shared-secret shared-secret;
        }
    }
```

```

pap-password pap-password;
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address ip-address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
user-group-profile profile-name;
}
domain-name-server;
domain-name-server-inet;
domain-name-server-inet6;
local {
    flat-file-profile profile-name;
}
preauthentication-order preauthentication-method;
provisioning-order (gx-plus | jsr | pcrf);
radius {
    accounting-server [ ip-address ];
    attributes {
        exclude {
            attribute-name packet-type;
            standard-attribute number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
                    | accounting-stop ];
            }
            vendor-id id-number {
                vendor-attribute vsa-number {
                    packet-type [ access-request | accounting-off | accounting-on | accounting-start
                        | accounting-stop ];
                }
            }
        }
    }
    ignore {
        dynamic-iflset-name;
        framed-ip-netmask;
        idle-timeout;
        input-filter;
        logical-system:routing-instance;
        output-filter;
        session-timeout;
        standard-attribute number;
        vendor-id id-number {
            vendor-attribute vsa-number;
        }
    }
}
authentication-server [ ip-address ];
options {

```

```
accounting-session-id-format (decimal | description);
calling-station-id-delimiter delimiter-character;
calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
}
chap-challenge-in-request-authenticator;
client-accounting-algorithm (direct | round-robin);
client-authentication-algorithm (direct | round-robin);
coa-dynamic-variable-validation;
ethernet-port-type-virtual;
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
juniper-dsl-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    pw-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
        adapter-width width;
        port-width width;
        slot-width width;
        vci-width width;
        vpi-width width;
    }
}
nas-port-id-delimiter delimiter-character;
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
```

```

    }
    nas-port-type {
        ethernet {
            port-type;
        }
    }
    revert-interval interval;
    service-activation {
        dynamic-profile (optional-at-login | required-at-login);
        extensible-service (optional-at-login | required-at-login);
    }
    vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    max-outstanding-requests value;
    source-address source-address;
    timeout seconds;
}
service {
    accounting {
        statistics (time | volume-time);
        update-interval minutes;
    }
    accounting-order (activation-protocol | local | radius);
}
session-options {
    client-idle-timeout minutes;
    client-idle-timeout-ingress-only;
    client-session-timeout minutes;
    strip-user-name {
        delimiter [ delimiter ];
        parse-direction (left-to-right | right-to-left);
    }
}
}
}

```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure PPP CHAP, or a profile and its subscriber access, L2TP, or PPP properties.

Options *profile-name*—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege admin—To view this statement in the configuration.
Level admin-control—To add this statement to the configuration.

**Related
Documentation**

- *Configuring the PPP Authentication Protocol*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- *Configuring L2TP Properties for a Client-Specific Profile*
- [Configuring an L2TP Access Profile on the LNS on page 246](#)
- [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring Service Accounting with JSRC*
- *Configuring Service Accounting in Local Flat Files*
- *AAA Service Framework Overview*

proxy-mode

Syntax	proxy-mode;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.</p> <p>You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>DHCP Relay Proxy Overview</i> • <i>Extended DHCP Relay Agent Overview</i> • <i>Enabling DHCP Relay Proxy Mode</i>

ps0 (Pseudowire Subscriber Interfaces)

Syntax	<pre>ps0 { anchor-point <i>lt-device</i>; mtu <i>bytes</i>; mac <i>mac-address</i>; no-gratuitous-arp-request; (flexible-vlan-tagging stacked-vlan-tagging untagged vlan-tagging); }</pre>
Hierarchy Level	[edit logical-systems transport-ls interfaces]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>Configure the pseudowire logical device.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Pseudowire Subscriber Logical Interfaces Overview on page 307• Configuring a Pseudowire Subscriber Logical Interface on page 313• Configuring a Pseudowire Subscriber Logical Interface Device on page 315• Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface on page 318• Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface on page 321


pseudowire-service (Pseudowire Subscriber Interfaces)

Syntax	<pre>pseudowire-service { device-count <i>number</i>; }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	<p>Configure properties for the pseudowire devices on the router.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Pseudowire Subscriber Logical Interfaces Overview on page 307• Configuring a Pseudowire Subscriber Logical Interface on page 313• Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 315

qualified-next-hop (Subscriber Management)

Syntax	<code>qualified-next-hop <i>interface-name</i> { <code>mac-address</code> <i>address</i>; }</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access-internal <code>route</code> <i>subscriber-ip-address</i>], [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access-internal <code>route</code> <i>subscriber-ip-address</i>], [edit dynamic-profiles <i>profile-name</i> routing-options <code>access-internal route</code> <i>subscriber-ip-address</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options route <i>subscriber-ip-address</i>] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> route <i>subscriber-ip-address</i>] hierarchy levels introduced in Junos OS Release 10.1.
Description	Dynamically configure the qualified next-hop and the MAC address for an access-internal route for DHCP and PPP subscriber interfaces.
Options	<i>interface-name</i> —Either the specific interface you want to assign to the access route or the variable, or the \$junos-interface-name variable. The variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in. The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32

reject-unauthorized-ipv6cp

Syntax	reject-unauthorized-ipv6cp;
Hierarchy Level	[edit protocols ppp-service]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Configure the router to reject any IPv6 Control Protocol (IPv6CP) negotiation messages on dynamic interfaces when no appropriate IPv6 address or prefix has been received from AAA. IPv6CP negotiation attempts are also rejected when only a Framed-IPv6-Prefix attribute is received but router advertisement is not enabled in the dynamic profile.
<div>  NOTE: IPv6CP negotiation messages are not rejected for static interfaces. </div>	
Default	IPv6CP negotiation is allowed regardless of the presence of IPv6 attributes received from AAA.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address</i>

relay-option-82

Syntax

```

relay-option-82 {
    circuit-id {
        include-irb-and-l2;
        keep-incoming-circuit-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    server-id-override
    vendor-specific{
        host-name;
        location;
    }
}

```

Hierarchy Level

```

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 forwarding-options dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group
 group-name]

```

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

To enable insertion of option 82 information in DHCP packets, you must specify at least one of the **circuit-id** or **remote-id** statements.

You can use the **relay-option-82** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level to control insertion of option 82 information globally, or at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level to control insertion of option 82 information for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the **delete relay-option-82** statement.

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

- *Using DHCP Relay Agent Option 82 Information*
- [dhcp-relay on page 431](#)

remote-gateway (Tunnel Profile)

Syntax `remote-gateway {
 address server-ip-address;
 gateway-name server-name;
}`

Hierarchy Level [edit access tunnel-profile *profile-name* [tunnel](#) *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address and hostname of the remote gateway at the L2TP tunnel endpoint, the LNS.


The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)

report-ingress-shaping-rate (Dynamic CoS Interfaces)

Syntax	<code>report-ingress-shaping-rate <i>bps</i>;</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 17.2 on MX Series.
Description	Report the ingress shaping rate in bits per second that is used as the receive speed (Rx) for the LAC to send to the LNS. The ingress shaping rate is used when the method for deriving the connect speed is configured as service-profile with the tx-connect-speed statement at the [edit services l2tp] hierarchy level.
	<div>  <p>NOTE: This statement is supported only when the effective shaping-rate statement is included at the [edit chassis] hierarchy level. If it is not, the subscriber login fails and a system log message is generated. There is no commit check failure because a commit check cannot be performed at run time.</p> </div>
Options	<p>bps—Ingress shaping rate in bits per second.</p> <p>Range: 1000 through 6,400,000,000,000 (6.4Tbps)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228 • Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219 • Guidelines for Configuring Dynamic CoS for Subscriber Access • Applying a Classifier to a Subscriber Interface in a Dynamic Profile

request services l2tp destination unlock


Syntax	<code>request services l2tp destination unlock <i>destination-name</i></code>
Release Information	Command introduced in Junos OS Release 13.2.
Description	Remove the specified destination from the destination lockout list immediately, before its lockout period expires. As a result, the L2TP process can again consider the destination during the selection of new tunnels.
Options	<i>destination-name</i> —Name of the destination to be unlocked.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Removing an L2TP Destination from the Destination Lockout List on page 165 • Configuring the L2TP Destination Lockout Timeout on page 165 • show services l2tp destination lockout on page 736
List of Sample Output	request services l2tp destination unlock on page 595
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output


request services l2tp destination unlock

```
user@host> request services l2tp destination unlock dest-a
Destination dest-a unlocked
```

retransmission-count-established (L2TP)

Syntax	<code>retransmission-count-established</code> <i>count</i> ;
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Set the maximum number of times control messages are retransmitted for established tunnels.
<div> BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing <code>no services l2tp tunnel retransmission-count-established</code>.</div>	
Options	<i>count</i> —Number of retransmissions. Range: 2 through 30 Default: 7
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Retransmission Attributes for L2TP Control Messages on page 171• Configuring an L2TP LAC on page 173• Configuring an L2TP LNS with Inline Service Interfaces on page 239

retransmission-count-not-established (L2TP)

Syntax	<code>retransmission-count-not-established <i>count</i>;</code>
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Set the maximum number of times control messages are retransmitted for tunnels that are not established.
<div>  <p>BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing <code>no services l2tp tunnel retransmission-count-not-established</code>.</p> </div>	
Options	<p><i>count</i>—Number of retransmissions.</p> <p>Range: 2 through 30</p> <p>Default: 5</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Retransmission Attributes for L2TP Control Messages on page 171 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239


route (Access)

Syntax	<pre>route <i>prefix</i> { next-hop <i>next-hop</i>; metric <i>route-cost</i>; preference <i>route-distance</i>; tag <i>route-tag</i>; }</pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options access]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access] hierarchy levels introduced in Junos OS Release 10.1.</p>
Description	Dynamically configure the parameters for access routes.
Options	<p><i>prefix</i>—Either the specific route prefix that you want to assign to the access route or one of the following route prefix variables.</p> <ul style="list-style-type: none">For IPv4 access routes, use the variable, \$junos-framed-route-ip-address-prefix. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].For IPv6 access routes, use the variable, \$junos-framed-route-ipv6-address-prefix. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99]. <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">Configuring Dynamic Access Routes for Subscriber Management on page 33

route (Access Internal)

Syntax	<pre>route <i>subscriber-ip-address</i> { next-hop <i>next-hop</i>; qualified-next-hop <i>underlying-interface</i> { mac-address <i>address</i>; } }</pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access-internal],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access-internal],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options access-internal]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access-internal] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access-internal] hierarchy levels introduced in Junos OS Release 10.1.</p>
Description	<p>Dynamically configure parameters for an access-internal route.</p>
Options	<p><i>subscriber-ip-address</i>—Either the specific IP address you want to assign to the access-internal route or the subscriber IP address variable (\$junos-subscriber-ip-address). The subscriber IP address variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32 • Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 123

route-suppression (DHCP Local Server and Relay Agent)

Syntax	<code>route-suppression (access access-internal destination);</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group group-name], [edit forwarding-options dhcp-relay dhcpv6 group group-name], [edit logical-systems <i>logical-system-name</i> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>...], [edit routing-instances <i>routing-instance-name</i> ...], [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server group group-name], [edit system services dhcp-local-server dhcpv6 group group-name]</p>
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Configure the <code>jdhcpd</code> process to suppress the installation of access, access-internal, or destination routes during client binding.
<div>  <p>NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.</p> </div>	
Options	<p>access—(DHCPv6 only) Suppress installation of access routes. You can use the access and access-internal options in the same statement for DHCPv6.</p> <p>access-internal—In a DHCPv4 hierarchy, suppress installation of both access-internal and destination routes. In a DHCPv6 hierarchy, suppress access-internal routes only. Can be configured in the same statement with the access option.</p> <p>destination—(DHCPv4 only) Suppress installation of destination routes. This option and the access-internal option are mutually exclusive; however, the access-internal option also suppresses destination routes.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 38

routing-instance (Tunnel Profile)

Syntax	<code>routing-instance <i>routing-instance-name</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify a routing instance for a tunnel.
Options	<i>routing-instance-name</i> —Name of the routing instance. Default: Routing instance <i>default</i>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Tunnel Profile for Subscriber Access on page 207

routing-instance (L2TP Destination)

Syntax	<pre>routing-instance <i>routing-instance-name</i> { drain; }</pre>
Hierarchy Level	[edit services l2tp destination address <i>ip-address</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the routing instance in which the destination is created.
Options	<i>routing-instance-name</i> — Name of the routing instance. Default: Routing instance <i>default</i> The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring L2TP Drain on page 166

routing-instance (L2TP Tunnel Destination)

Syntax	<code>routing-instance <i>routing-instance-name</i> { drain; }</code>
Hierarchy Level	[edit services l2tp tunnel name <i>tunnel-name</i> address <i>ip-address</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify the routing instance in which the tunnel to the destination address is created.
Options	<p><i>routing-instance-name</i>— Name of the routing instance.</p> <p>Default: Routing instance <i>default</i></p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Drain on page 166

routing-instances (Dynamic Profiles)

```

Syntax  routing-instances routing-instance-name {
        interface interface-name;
        multicast-snooping-options {
        }
        routing-options {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
                    preference route-distance;
                    tag route-tag;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
        rib routing-table-name {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
                    preference route-distance;
                    tag route-tag;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
    }
}

```

Hierarchy Level [edit dynamic-profiles]
[edit logical-systems *logical-system-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Support at the **logical-systems** hierarchy level was introduced in Junos OS Release 14.2.

Description Dynamically configure an additional routing entity for a router.

Options *routing-instance-name*—The routing instance variable (*\$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution*

routing-options (Dynamic Profiles)

Syntax	<pre> routing-options { access { route prefix { metric route-cost; next-hop next-hop; preference route-distance; tag route-tag; } } access-internal { route subscriber-ip-address { qualified-next-hop underlying-interface { mac-address address; } } } multicast { interface interface-name { no-qos-adjust; } } rib routing-table-name { access { route prefix { metric route-cost; next-hop next-hop; preference route-distance; tag route-tag; } } access-internal { route subscriber-ip-address { qualified-next-hop underlying-interface { mac-address address; } } } } } </pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i>], [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6. Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance] hierarchy level introduced in Junos OS Release 10.1.</p>
Description	<p>Configure protocol-independent routing properties in a dynamic profile.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Dynamic Access Routes for Subscriber Management on page 33](#)
- [Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 32](#)

rule (IP Reassembly)

Syntax

```
rule rule-name {  
    match-direction direction;  
}
```

Hierarchy Level [edit services [ip-reassembly](#)]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure an IP reassembly rule, which is used for inline IP reassembly on the inline services interface. The IP reassembly rule can be attached to a service set to indicate that the service set is of type IP reassembly. For inline IP reassembly, each rule must include the **match-direction** statement, which specifies the direction in which the match is applied.



NOTE: If you configure an IP reassembly rule, then you must configure the **match-direction** statement.

Options **rule-name**—Name of the IP reassembly rule.
Range: Up to 63 characters

The remaining statement is explained separately. See [CLI Explorer](#).



NOTE: To create more than one IP reassembly rule, include the rule statement multiple times.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring IP Inline Reassembly for L2TP on page 290](#)
- [IP Packet Fragment Reassembly for L2TP Overview on page 289](#)

rx-connect-speed-when-equal (L2TP LAC)

Syntax	rx-connect-speed-when-equal
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Enable sending the receive connect speed, which is represented by AVP 38, even when its value is equal to that of the transmit connect speed, which is represented by AVP 24. By default, AVP 38 is sent from the LAC to the LNS during the establishment of an L2TP tunnel session, only when its value is different from AVP 24. You can override the default setting with this configuration statement.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228• Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 227

rx-window-size (L2TP)

Syntax	<code>rx-window-size packets;</code>
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Configure the L2TP receive window size for an L2TP tunnel.
Options	<p><i>packets</i>—Number of packets that a peer can transmit without receiving an acknowledgment from the router. By default, this value is set to 4 packets. If the receive window size is configured to its default value, the router does not send the Receive Window Size AVP (AVP 10) in the first tunnel negotiation packet that is sent to its peer.</p> <p>Range: 4 through 128</p> <p>Default: 4</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Setting the L2TP Receive Window Size on page 163• Configuring an L2TP LAC on page 173• Configuring an L2TP LNS with Inline Service Interfaces on page 239

secondary-interface (Aggregated Inline Services)

Syntax	<code>secondary-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit interfaces asix aggregated-inline-services-options]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Specify the secondary (backup) inline services member link in the asi bundle. You must also configure a primary (active) member link on a different MPC with the primary-interface statement. The secondary member provides 1:1 redundancy for subscriber service sessions on the primary member link. The bandwidth configured at the [edit chassis fpc slot pic number inline-services bandwidth] hierarchy level must be the same for both member links.</p> <p>Redundancy is not achievable if you configure the primary and secondary interface on the same MPC, because both member interfaces go down if the card goes down. Consequently, if you configure both interfaces on the same MPC, the subsequent configuration commit fails.</p>
Options	<i>interface-name</i> —Name of an inline services physical interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces on page 255 • Configuring an L2TP LNS with Inline Service Interfaces on page 239

secret (Tunnel Profile)

Syntax	<code>secret password;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the tunnel password sent to the LNS for authentication.
Options	<i>password</i> —Cleartext password.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

service-device-pool (L2TP)

Syntax	<code>service-device-pool pool-name;</code>
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Assign a pool of service interfaces to the tunnel group to balance traffic across.




NOTE: 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.

Options	<i>pool-name</i> —Name of the service device pool.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274

service-device-pools (L2TP Service Interfaces)

Syntax	<pre> service-device-pools { pool pool-name { interface service-interface-name; } } </pre>
Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Configure one or more pools of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.
Options	<p>pool-name—Name of the service interface pool.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 284

service-interface (L2TP Processing)

Syntax	<code>service-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. <i>si-fpc/pic/port</i> option added in Junos OS Release 11.4. Option <i>asifpc</i> added in Junos OS Release 16.2.
Description	Specify the service interface responsible for handling L2TP processing.
	<div> NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.</div>
Options	<p><i>interface-name</i>—Name of the service interface. The ae, si, and sp interface types are supported as follows:</p> <ul style="list-style-type: none">• <i>asix</i>—(MPCs on MX Series routers) Aggregated inline services interface.• <i>sp-fpc/pic/port</i>—On AS or Multiservices PICs on M7i, M10i, and M120 routers.• <i>si-fpc/pic/port</i>—On MPCs on MX Series routers.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Local Gateway Address and PIC• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274

service-profile (L2TP)

Syntax	<code>service-profile <i>profile-name</i>(<i>parameter</i>)&<i>profile-name</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> <i>client</i> <i>client-name</i> <i>l2tp</i>]</code> <code>[edit services <i>l2tp</i> tunnel-group <i>group-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 18.1R1 on MX Series routers.
Description	<p>Configure one or more dynamic service profiles to be applied to subscriber sessions at activation for all subscribers in the specified tunnel group or on the specified LAC. Services are typically applied to L2TP sessions with RADIUS VSAs or CoA requests. In multivendor environments, you might use only standard attributes to simplify management of multiple vendor VSAs. This statement enables you to apply services without using an external authority such as RADIUS. The locally configured list of services (service profiles) serves as local authorization that is applied by authd during client session activation. This list of services is subject to the same validation and processing as services originating from an external authority, such as RADIUS.</p> <p>You can optionally specify parameters that are passed to the corresponding service when it is activated for the session. The parameter might override values configured in the profile itself, such as a downstream shaping rate for a CoS service. This enables you to use the same service profile for multiple situations with different requirements, or to modify a previously applied value for a service.</p> <p>You can still use RADIUS VSAs or CoA requests together with the service profiles. If services are sourced from an external authority as authorization during authentication or during subscriber session provisioning (activation), the services from the external authority take strict priority over those in the local configuration. If a service applied with RADIUS is the same as a service applied with a service profile in the CLI, but with different parameters, the RADIUS service is applied with a new session ID and takes precedence over the earlier service profile.</p> <p>When service profiles are configured on a LAC client and on a tunnel group that uses that LAC client, the LAC configuration overrides the tunnel group configuration. Only the service profile configured on the LAC client is applied to subscribers in the tunnel group.</p>
Options	<p><i>profile-name</i>—Name of a dynamic service profile that defines a service to be applied to L2TP subscriber sessions. You can specify one or more service profiles, separated by an ampersand (&).</p> <p><i>parameter</i>—(Optional) Value to be passed to the service when it is activated on the subscriber session.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

**Related
Documentation**

- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)
- [Configuring an L2TP Access Profile on the LNS on page 246](#)

service-rate-limiter (Access)

Syntax	<pre>service-rate-limiter { rx-multiplier <i>rx-multiplier</i>; service-name <i>service-profile-name</i>; tx-multiplier <i>tx-multiplier</i>; }</pre>
Hierarchy Level	[edit access]
Release Information	Statement introduced in Junos OS Release 18.1R1 on MX Series routers.
Description	<p>Specify a dynamic service profile that provides rates for upstream and downstream traffic that the LAC communicates to the LNS. When the Juniper Networks Activate-Service VSA (26-65) is received in the RADIUS Access-Accept message at subscriber login, the VSA is evaluated to determine whether the configured name is also conveyed in the VSA. If it is, the rate values are collected and stored in the session database for the subscriber and then sent in the ICCN message to the LNS. You can either define the rate values as defaults in the service profile or configure them to be passed by RADIUS in VSA 26-65. When they are passed by the VSA, the first value is taken as the receive speed (the upstream rate from the subscriber to the LAC) and the second value is taken as the transmit speed (the downstream rate from the LAC to the subscriber).</p> <p>The multipliers convert the rates from Kbps to bps, which is required for the AVPs. You can also use the multiplier options to adjust the rates up or down. The adjusted values correspond to the Juniper Networks RADIUS VSAs, Rx-Connect-Speed (26-163) and Tx-Connect-Speed (26-162). These values are stored as such in the session database (SDB). Because the values are available in the SDB before the L2TP connection is negotiated, the LAC includes them in the ICCN message as AVP 38 (Rx connect speed) and AVP 24 (Tx connect speed). The rate values are treated as RADIUS-sourced values and consequently have the highest precedence among multiple sources.</p>
Options	<p><i>rx-multiplier</i>—(Optional) Multiplier applied to convert the Rx connect speed value Kbps to bps and optionally adjust the rate up or down.</p> <p>Default: 1000</p> <p>Range: 1 through 2000</p> <p><i>service-profile-name</i>—Name of the dynamic service profile conveyed in VSA 26-65 that specifies upstream and downstream traffic rates.</p> <p><i>tx-multiplier</i>—(Optional) Multiplier applied to convert the Tx connect speed value Kbps to bps and optionally adjust the rate up or down.</p> <p>Default: 1000</p> <p>Range: 1 through 2000</p>
Required Privilege Level	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds on page 235](#)
 - [Configuring an L2TP LAC on page 173](#)
 - [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)

session-mode

Syntax	session-mode (automatic multihop singlehop);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the session mode.
Options	<p>Default: automatic</p> <p>automatic—Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.</p> <p>multihop—Configure multihop BFD sessions and passive DHCP clients.</p> <p>single-hop—Configure single hop BFD sessions and non-passive DHCP clients.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94• Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

session-options

Syntax	<pre> session-options { client-group [<i>group-names</i>]; client-idle-timeout <i>minutes</i>; client-idle-timeout-ingress-only; client-session-timeout <i>minutes</i>; strip-user-name { delimiter [<i>delimiter</i>]; parse-direction (left-to-right right-to-left); } } </pre>
Hierarchy Level	[edit access profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	<p>(MX Series and SRX Series devices) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.</p> <p>(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Session Options for Subscriber Access</i> • <i>Configuring Subscriber Session Timeout Options</i> • Configuring Username Modification for Subscriber Sessions on page 179 • <i>Removing Inactive Dynamic Subscriber VLANs</i>

sessions-limit-group (L2TP)

Syntax	<code>sessions-limit-group <i>limit-group-name</i> { maximum-sessions <i>number</i>; }</code>
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 16.1.
Description	Create a group of clients so that you can limit the number of L2TP sessions allowed for the client group. You can create up to 200 groups.
Options	<i>limit-group-name</i> —Identifier of the session-limit group for which the session limit is configured. The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204• Configuring an L2TP LAC on page 173• Configuring an L2TP LNS with Inline Service Interfaces on page 239• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274• L2TP for Subscriber Access Overview on page 147

sessions-limit-group (L2TP Client Profile)

Syntax	<code>sessions limit-group <i>limit-group-name</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]</code>
Release Information	Statement introduced in Junos OS Release 16.1.
Description	Specify in an L2TP access profile the session limit group to which a client is assigned by the profile.
Options	<i>limit-group-name</i> —Identifier of the session-limit group to which a client is assigned.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Limiting the Number of L2TP Sessions Allowed by the LAC or LNS on page 204 • Configuring an L2TP LAC on page 173 • Configuring an L2TP LNS with Inline Service Interfaces on page 239 • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274 • L2TP for Subscriber Access Overview on page 147

shared-secret

Syntax	<code>shared-secret <i>shared-secret</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i> client <i>client-name</i> l2tp]</code>
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Configure the shared secret.
Options	<i>shared-secret</i> —Shared secret key for authenticating the peer.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring L2TP Properties for a Client-Specific Profile • Configuring an L2TP Access Profile on the LNS on page 246

soft-gre

Syntax

```

soft-gre {
    destination-networks {
        [prefix];
    }
    dynamic-profile profile-name;
    group-name;
    service-interface psn;
    source-address wag-ip-address;
    <tunnel-idle-timeout seconds>;
}

```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 17.2R1 on MX Series routers.

Description Configure the conditions for enabling dynamic-bridged generic routing encapsulation (GRE) tunnel creation on the MX Series router Wi-Fi access gateway (WAG).



NOTE: Configuration of multiple dynamic tunnel groups is supported.

Options **destination-networks [prefix]**—Use the specified IP subnets from which soft-GRE connection requests from the customer can be processed.

group-name—Name of the dynamic GRE tunnel group.

profile-name—Name of the dynamic profile that creates the tunnel.



NOTE: To support VLAN autosensing on a GRE tunnel, you must also specify the auto-configure options at the [edit dynamic-profile *profile-name* interfaces unit] hierarchy level. These options include a reference to the dynamic profile that creates VLANs.

service-interface psn—Use the specified pseudowire subscriber interface device (IFD) on which the tunnels are built.

source-address wag-ip-address—Use the specified source IP address of the GRE tunnels for the WAG. This is the IP address on which incoming GRE traffic must be received by the MX Series router.

tunnel-idle-timeout seconds—(Optional) Use the specified number of seconds that a GRE tunnel remains up after the last subscriber session on the tunnel has ended. If

set to 0, then no idle timeout is set, and the tunnel remains up for an unlimited period of time.

Range: 0 through 65535

Default: 120

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Wi-Fi Access Gateway Deployment Model Overview on page 344](#)
- [Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345](#)
- [Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation on page 349](#)
- [show services soft-gre tunnel on page 779](#)

source-gateway (Tunnel Profile)

Syntax

```
source-gateway {
  address client-ip-address;
  gateway-name client-name;
}
```

Hierarchy Level [edit access tunnel-profile *profile-name* **tunnel** *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address and hostname of the source gateway at the local L2TP tunnel endpoint, the LAC.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)

stacked-vlan-tagging

Syntax	stacked-vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface. For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i>


statistics (Access Profile)

Syntax	statistics (time volume-time);
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. volume-time option added in Junos OS Release 9.4.
Description	Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.
Options	time —Collect uptime statistics only. volume-time —Collect both volume and uptime statistics.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i>

strip-user-name (Access Profile)

Syntax	<pre>strip-user-name { delimiter <i>delimiter</i>; parse-direction (left-to-right right-to-left); }</pre>
Hierarchy Level	[edit access profile <i>profile-name</i> session-options]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	<p>Configure the details of username stripping for a subscriber access profile. This configuration determines how a portion of a subscriber login string is identified and discarded, leaving the remainder for subsequent use as a modified username by an external AAA server for session authentication and accounting. For example, the modified username might appear in RADIUS Access-Request, Acct-Start, and Acct-Stop messages, as well as RADIUS-initiated disconnect requests and change of authorization (CoA) requests.</p> <p>You can use the aaa-options <i>aaa-options-name</i> statement at the [edit access] hierarchy level to define options that specify the LS:RI context for AAA authorization and configuration for a specific subscriber or a set of subscribers, overriding any such configuration within a domain map.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Session Options for Subscriber Access • Configuring Username Modification for Subscriber Sessions on page 179 • Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245 • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

subscriber-context (AAA Options)

Syntax	subscriber-context <i>subscriber-context-name</i> ;
Hierarchy Level	[edit access aaa-options <i>aaa-options-name</i>]
Release Information	Statement introduced in Junos OS Release 16.2.
Description	Specify the logical-system:routing-instance (LS:RI) in which the subscriber interface is placed. For example, this may correspond to the LAC-facing interface on the LNS that is accessed by all requests from a subscriber residence.
<div> NOTE: Only the default (master) logical system is supported.</div>	
Options	<i>subscriber-context-name</i> —Name of the logical-system:routing-instance in which the subscriber interface is placed.
Required Privilege Level	access—To view this statement in the configuration. access-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Session Options for Subscriber Access• Configuring Username Modification for Subscriber Sessions on page 179• Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 245• Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface on page 242

subscriber-management (Subscriber Management)

```
Syntax subscriber-management {
    enable;
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
}
overrides {
    interfaces {
        family (inet | inet6) {
            layer2-liveness-detection;
        }
    }
    no-unsolicited-ra;
    shmlog {
        disable;
        file filename <files maximum-no-files> <size maximum-file-size>;
        filtering enable;
        log-name {
            all;
            logname {
                <brief | detail | extensive | none | terse>;
                <file-logging | no-file-logging>;
            }
        }
        log-type (debug | info | notice);
    }
    |
}
traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
}
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation


- *Configuring the Router to Strictly Enforce the Subscriber Scaling License*

- [Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 108](#)
- [Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events](#)
- [Tracing Subscriber Management Database Operations for Subscriber Access](#)
- [Configuring Junos OS Enhanced Subscriber Management](#)
- [DHCP Liveness Detection Using ARP and Neighbor Discovery Packets on page 98](#)


tag (Access)

Syntax	<code>tag route-tag;</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> , <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> , <code>[edit dynamic-profiles <i>profile-name</i> routing-options access route <i>prefix</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.2. Support at the <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> and <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> hierarchy levels introduced in Junos OS Release 10.1.
Description	Dynamically configure the tag for an access route.
Options	route-tag —Either the specific tag you want to assign to the access route or either of the following tag variables: <ul style="list-style-type: none">• \$junos-framed-route-tag—Tag assigned to an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-Route attribute [22].• \$junos-framed-route-ipv6-tag—Tag assigned to an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-IPv6-Route attribute [99].
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Configuring Dynamic Access Routes for Subscriber Management on page 33

threshold (detection-time)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd detection-time], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd detection-time], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div>  NOTE: The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval. </div>	
Options	<p><i>milliseconds</i>— Value for the detection time adaptation threshold.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

threshold (transmit-interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval]</p>
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
Options	<p><i>milliseconds</i> — Threshold value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p>
<div>  <p>NOTE: The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89

tos-reflect (L2TP)

Syntax	tos-reflect;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header. When CoS rewrite rules are also configured ([edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]), the rewrite is performed only on the inner IP ToS; the outer IP TOS value is not affected.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Dynamic CoS for an L2TP LNS Inline Service</i>

trace (DHCP Relay Agent)

Syntax	trace;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>EX Series switches do not support DHCPv6.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i>• <i>Tracing Extended DHCP Operations</i>• <i>Tracing Extended DHCP Operations for Specific Interfaces</i>

traceoptions (Services L2TP)

Syntax

```

traceoptions {
  debug-level level;
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    protocol name;
    user user@domain;
    user-name username;
  }
  flag flag;
  interfaces interface-name {
    debug-level level;
    flag flag;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}

```

Hierarchy Level [edit services [l2tp](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define tracing operations for L2TP processes.

Options **debug-level *level***—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:

- **detail**—Trace detailed debug information.
- **error**—Trace error information.
- **packet-dump**—Trace packet decoding information.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

- **protocol *name***—One of the following protocols; this option does not apply to L2TP on MX Series routers:

- **l2tp**
- **ppp**
- **radius**
- **udp**
- **user** *user@domain*—Username of a subscriber; this option does not apply to L2TP on M Series routers. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.
- **user-name** *username*—Username of a subscriber; this option does not apply to L2TP on MX Series routers.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **configuration**—Trace configuration events.
- **events**—Trace interface events.
- **general**—Trace general events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization.
- **ipc-rx**—Trace IPC receive events.
- **ipc-tx**—Trace IPC transmit events.
- **memory**—Trace memory management code.
- **message**—Trace message processing code.
- **packet-error**—Trace packet error events.
- **parse**—Trace parsing events.
- **protocol**—Trace L2TP events.
- **receive-packets**—Trace received L2TP packets.
- **routing-process**—Trace routing process interactions.
- **routing-socket**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **states**—Trace state machine events.
- **timer**—Trace timer events.
- **transmit-packets**—Trace transmitted L2TP packets.
- **tunnel**—Trace tunnel events.

interfaces *interface-name*—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- **debug-level *level***—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
 - **detail**—Trace detailed debug information.
 - **error**—Trace error information.
 - **extensive**—Trace all PIC debug information.
- **flag *flag***—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:
 - **all**—Trace everything.
 - **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
 - **packet-dump**—Dump each packet content based on debug level.
 - **protocol**—Trace L2TP, PPP, and multilink handling.
 - **system**—Trace packet processing on the PIC.

level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege	trace—To view this statement in the configuration.
Level	trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Tracing L2TP Operations</i>• Tracing L2TP Operations for Subscriber Access on page 371
------------------------------	--

traceoptions (Protocols PPP Service)

Syntax

```

traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    filter {
        aci regular-expression;
        ari regular-expression;
        service-name regular-expresion;
        underlying-interface interface-name;
        user user@domain;
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}

```

Hierarchy Level [edit protocols ppp-service]

Release Information Statement introduced in Junos OS Release 9.5.
user option added in Junos OS Release 14.1.

Description Define tracing operations for PPP service processes.

Options **file filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files number—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.



BEST PRACTICE: Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- **aci regular-expression**—Regular expression to match the agent circuit identifier provided by PPP client.

- **ari *regular-expression***—Regular expression to match the agent remote identifier provided by PPP client.
- **service *regular-expression***—Regular expression to match the name of PPPoE service.
- **underlying-interface *interface-name***—Name of a PPP underlying interface. You cannot use a regular expression for this filter option.
- **user *user@domain***—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **accounting-statistics**—Trace accounting statistics events.
- **all**—Trace all operations.
- **authentication**—Trace authentication events.
- **chap**—Trace CHAP events.
- **events**—Trace interface events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization events.
- **interface-db**—Trace interface database events.
- **lcp**—Trace LCP state machine events.
- **memory**—Trace memory processing events.
- **ncp**—Trace NCP state machine events.
- **packet-error**—Trace packet error events.
- **pap**—Trace PAP events.
- **parse**—Trace parsing events.
- **profile**—Trace libdynamic profile events.
- **receive-packets**—Trace received PPP packets.
- **routing-process**—Trace routing process interactions.
- **rtp**—Trace real-time priority events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **smi-services-sentry**—Trace SMI services requests and retries.
- **states**—Trace state machine events.
- **transmit-packets**—Trace transmitted PPP packets.
- **tunnel**—Trace L2TP tunneling events.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **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

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level	<p>trace—To view this statement in the configuration.</p> <p>trace-control—To add this statement to the configuration.</p>
---------------------------------	--

Related Documentation	<ul style="list-style-type: none"> • Tracing PPP Service Operations for Subscriber Access on page 363
------------------------------	--

tracoptions (Subscriber Management)

Syntax	<pre>tracoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; }</pre>
Hierarchy Level	[edit system services subscriber-management]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Define tracing operations for subscriber management interface processes.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none">• all—Trace all operations.• database—Trace database events.• general—Trace general events.• issu—Trace unified ISSU events.• server—Trace server events.• session-db—Trace session database interactions.• ui—Trace user interface events. <p>match <i>regular-expression</i>—(Optional) Refine the output to include lines that contain the regular expression.</p> <p>no-world-readable—(Optional) Disable unrestricted file access.</p> <p>size <i>maximum-file-size</i>—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.</p> <p>Syntax: sizek to specify KB, sizem to specify MB, or sizeg to specify GB</p> <p>Range: 10240 through 1073741824</p>

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- *Tracing Subscriber Management Database Operations for Subscriber Access*

transmit-interval

Syntax transmit-interval {
 threshold *milliseconds*;
 minimum-interval *milliseconds*;
}

Hierarchy Level [edit system services dhcp-local-server liveness-detection method [bfd](#)],
[edit system services dhcp-local-server dhcpv6 liveness-detection method [bfd](#)],
[edit forwarding-options dhcp-relay liveness-detection method [bfd](#)], [edit forwarding-options
 dhcp-relay dhcpv6 liveness-detection method [bfd](#)],
[edit system services dhcp-local-server group *group-name* liveness-detection method [bfd](#)],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method
 [bfd](#)],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method [bfd](#)],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method
 [bfd](#)]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the Bidirectional Forwarding Detection (BFD) transmit interval.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94](#)
- [Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89](#)

tunnel (L2TP)

Syntax tunnel {
 assignment-id-format (assignment-id | client-server-id);
 idle-timeout *seconds*;
 maximum-sessions *number*;
 minimum-retransmission-timeout;
 name *name* {
 address *ip-address* {
 drain;
 routing-instance *routing-instance-name* {
 drain;
 }
 }
 drain;
 }
 nas-port-method;
 retransmission-count-established *count*;
 retransmission-count-not-established *count*;
 rx-window-size *packets*;
 tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port | reject |
 reject-ip-address | reject-udp-port);
 }

Hierarchy Level [edit services [l2tp](#)]

Release Information Statement introduced in Junos OS Release 11.4.
 rx-window-size option added in Junos OS Release 13.2.

Description Configure L2TP tunnel characteristics.

 The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring an L2TP LAC on page 173](#)
 • [Configuring an L2TP LNS with Inline Service Interfaces on page 239](#)

tunnel (Tunnel Profile)

Syntax `tunnel tunnel-id {
 identification name;
 logical-system logical-system-name;
 max-sessions number;
 medium type;
 nas-port-method cisco-avp;
 preference number;
 remote-gateway {
 address server-ip-address;
 gateway-name server-name;
 }
 routing-instance routing-instance-name;
 secret password;
 source-gateway {
 address client-ip-address;
 gateway-name client-name;
 }
 type tunnel-type;
 }`

Hierarchy Level [edit access [tunnel-profile profile-name](#)]

Release Information Statement introduced in Junos OS Release 10.4.

Description Define the attributes of a tunnel for the tunnel profile. You can define up to 31 tunnels for each tunnel profile.

Options *tunnel-id*—Unique integer that identifies a tunnel defined within a profile. For a subscriber, RADIUS attributes and VSAs can supply or override the attributes defined here for the tunnel.

Range: 1 through 31

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring a Tunnel Profile for Subscriber Access on page 207](#)

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;`
 `maximum-sessions number;`
 `ppp-access-profile profile-name;`
 `receive-window packets;`
 `retransmit-interval seconds;`
 `service-device-pool pool-name;`
 `service-interface interface-name;`
 `service-profile profile-name(parameter)&profile-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.



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. See [CLI Explorer](#).

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](#)
- [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 274](#)

tunnel-profile (L2TP Tunnel Switching)

Syntax tunnel-profile *profile-name*;

Hierarchy Level [edit access [tunnel-switch-profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the name of an L2TP tunnel profile that defines the tunnel to which PPP subscriber traffic is switched.

Options *profile-name*—Unique name that identifies the tunnel profile; configured with the **tunnel-profile** statement at the **[edit access]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring L2TP Tunnel Switching on page 161](#)

tunnel-profile (Tunnel Profile)

Syntax	<pre>tunnel-profile <i>profile-name</i> { tunnel <i>tunnel-id</i> { identification <i>name</i>; logical-system <i>logical-system-name</i>; max-sessions <i>number</i>; medium <i>type</i>; nas-port-method cisco-avp; preference <i>number</i>; remote-gateway { address <i>server-ip-address</i>; gateway-name <i>server-name</i>; } routing-instance <i>routing-instance-name</i>; secret <i>password</i>; source-gateway { address <i>client-ip-address</i>; gateway-name <i>client-name</i>; } type <i>tunnel-type</i>; } }</pre>
Hierarchy Level	[edit access]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Define a tunnel profile for subscriber access.
Options	<p><i>profile-name</i>—Unique name that identifies the tunnel profile. The profile can be referenced from within a domain map or by the RADIUS Tunnel-Group VSA [26-64].</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

tunnel-switch-profile (L2TP Tunnel Switching, Application)

Syntax	<code>tunnel-switch-profile <i>profile-name</i>;</code>
Hierarchy Level	[edit access domain map <i>domain-map-name</i>], [edit services l2tp], [edit services l2tp <i>tunnel-group</i> <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Specify a tunnel switch profile that determines whether packets in an L2TP session from a LAC are switched to another session that has a different destination LNS.
Options	<i>profile-name</i> —Unique name that identifies the tunnel switch profile.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring L2TP Tunnel Switching on page 161

tunnel-switch-profile (L2TP Tunnel Switching, Definition)

Syntax tunnel-switch-profile *profile-name* {
 avp {
 bearer-type *action*;
 calling-number *action*;
 cisco-nas-port-info *action*;
 }
 tunnel-profile *profile-name*;
 }

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 13.2.

Description Define a tunnel switch profile for subscriber access; specify actions to take for L2TP AVPs in the switched packets and the profile for the tunnel to which the PPP traffic is switched.

The remaining statements are explained separately. See [CLI Explorer](#).

Options *profile-name*—Unique name that identifies the tunnel switch profile. The profile can be applied as a default or referenced from within a domain map, a tunnel group, or by the RADIUS Tunnel-Group VSA [26-64].

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation • [Configuring L2TP Tunnel Switching on page 161](#)

tx-address-change (L2TP LAC)

Syntax	tx-address-change (accept ignore ignore-ip-address ignore-udp-port reject reject-ip-address reject-udp-port);
Hierarchy Level	[edit services l2tp tunnel]
Release Information	Statement introduced in Junos OS Release 13.2. reject , reject-ip-address , and reject-udp-port options added in Junos OS Release 16.1.
Description	<p>Configure whether the LAC accepts, ignores, or rejects requests from an LNS to change the destination IP address, UDP port, or both.</p> <p>When configured to ignore change requests, the LAC continues to send packets to the original address or port, but accepts packets from the new address or port.</p> <p>When configured to reject change requests, the LAC sends a StopCCN message to the original address or port and then terminates the connection to that LNS. This method has precedence over the ignore configuration.</p>
Default	The LAC accepts IP address or UDP port change requests from the LNS.
Options	<p>accept—Accept all change requests for IP address or UDP port.</p> <p>ignore—Ignore all change requests for IP address or UDP port.</p> <p>ignore-ip-address—Ignore a change request for IP address, but accept a change request for UDP port.</p> <p>ignore-udp-port—Ignore a change request for UDP port, but accept a change request for IP address.</p> <p>reject—Reject all change requests for IP address or UDP port.</p> <p>reject-ip-address—Reject a change request for IP address, but accept a change request for UDP port.</p> <p>reject-udp-port—Reject a change request for UDP port, but accept a change request for IP address.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring How the LAC Responds to Address and Port Changes Requested by the LNS on page 175

tx-connect-speed-method (L2TP LAC)

Syntax tx-connect-speed-method *method*;

Hierarchy Level [edit services [l2tp](#)]

Release Information Statement introduced in Junos OS Release 11.4.
Options **ancp**, **pppoe-ia-tag**, and **static** introduced in Junos OS Release 13.1.
Option **static** deprecated in Junos OS Release 15.1.
Options **actual** and **none** added in Junos OS Release 15.1.
Option **actual** deprecated in Junos OS Release 17.2.
Option **service-profile** added in Junos OS Release 17.2.
Option **static** undeprecated in Junos OS Release 17.2.

Description Specify the method that determines how to derive the connect speed values sent from the LAC to the LNS.

When the session is being established, the speeds are included in the Incoming-Call-Connected (ICCN) message. The transmit speed is conveyed in AVP 24 (Tx-Connect-Speed) and the receive speed is conveyed in AVP 38 (Rx-Connect-Speed). Both values are in bits per seconds (bps). The LAC typically uses the **static** or **pppoe-ia-tags** method, because values for other configured methods are not available when the session is being established.

When connect speed updates are configured, the LAC sends the updated values for each session to the LNS in Connect-Speed-Update-Notification (CSUN) messages. The updated speeds are conveyed in the Connect Speed Update AVP (97).

When connection speed values are not available from the configured method, the LAC falls back to another source for the values. See [“Transmission of Tx and Rx Connection Speeds from LAC to LNS” on page 219](#) for tables describing the LAC fallback behavior by Junos OS release.

Options *method*—Method used to derive the connection speed values.

- **actual**—(Junos OS Releases 15.1, 16.1, 16.2, 17.1) The speed is derived from the CoS effective shaping rate that is enforced on the level 3 node based on local policy. In the supported releases, **actual** is the default method and has the highest preference among all configured methods.

This method is not available starting in Junos OS Release 17.2. However, it is configurable in the Tunnel-Tx-Speed-Method VSA (26-94). If you do so, it is translated to the **service-profile** method.

- **ancp**—The speed is derived from the configured ANCP value for the underlying interface. This value results from a user-defined percentage correction to the values received from the access node; this is configured per subscriber access line. The percentage accounts for encapsulation differences between, the router, the access loop, and the Layer 1 transport overhead. The initial rate sent to the LNS is

the ANCP value reported at the time the ICCN is sent. The ANCP value is not available for the ICCN message and falls over to another method. You can change the configured correction after a subscriber has logged in, but those changes do not affect the actual rate used by the LNS for that subscriber.

- **none**—This method prevents the LAC from sending AVP 24 and AVP 38 to the LNS. This option also overrides the Juniper Networks RADIUS VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163).
- **pppoe-ia-tags**—The speed is derived from the PPPoE IA tags received by the LAC from the DSLAM. This speed value is transmitted when a subscriber logs in and it cannot subsequently be changed. The value of Actual-Data-Rate-Downstream (VSA 26-129) is used for AVP 24. The value of Actual-Data-Rate-Upstream (VSA 26-130) is used for AVP 38; it is sent only when the values differ.



NOTE: This speed derived from the IA tags does not apply to subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved.

- **service-profile**—(Junos OS Releases 17.2 and higher) The downstream (Tx) speed is derived from the actual CoS that is enforced on the L3 node based on local policy. The upstream (Rx) speed is the value configured in the dynamic service profile; no adjustment is made to this value. The service-profile value is not available for the ICCN message and falls over to another method.

The **service-profile** method is supported only when the **effective shaping-rate** statement is included at the **[edit chassis]** hierarchy level. If it is not, the commit check fails. If the method is received from RADIUS in VSA 26-94, a system log message is generated instead, because no commit check is performed in this case.

- **static**—(Junos OS Releases 13.3, 14.1, and 14.2; Junos OS Releases 17.2 and higher) The speed is derived from the configured static Layer 2 speed. For Ethernet VLANs, this is the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory shaping rate is not configured on the underlying interface, then the actual speed of the underlying physical port is used. In the supported releases, **static** is the default method.

In Junos OS Releases 15.1, 16.1, 16.2, and 17.2, the **static** method is configurable for backward compatibility, but it is not supported. If you configure this method in the CLI or in the Tunnel-Tx-Speed-Method VSA (26-94), the LAC falls back to the port speed of the subscriber access interface.



NOTE: For ge and xe interfaces, the port speed value is set to 1,000,000,000 and for ae interfaces, the port speed value is set to 0. The value is sent in both AVP 24 and AVP 38.

Default:

- **static** (Starting in Junos OS Release 17.2)
- **actual** (Junos OS Releases 15.1, 16.1, 16.2, 17.1)
- **static** (Junos OS Releases 13.3, 14.1, and 14.2)

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS on page 228](#)
- [Transmission of Tx and Rx Connection Speeds from LAC to LNS on page 219](#)

type (Tunnel Profile)

Syntax	<code>type <i>tunnel-type</i>;</code>
Hierarchy Level	[edit access tunnel-profile <i>profile-name</i> tunnel <i>tunnel-id</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the tunnel type (protocol).
Default	l2tp
Options	<i>tunnel-type</i> —Tunnel protocol type. The only value currently available is l2tp .
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Tunnel Profile for Subscriber Access on page 207

unit (Dynamic PPPoE)

```

Syntax  unit logical-unit-number {
        keepalives interval seconds;
        no-keepalives;
        pppoe-options {
            underlying-interface interface-name;
            server;
        }
        ppp-options {
            aaa-options aaa-options-name;
            authentication [ authentication-protocols ];
            mru size;
            mtu (size | use-lower-layer);
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            ignore-magic-number-mismatch;
            initiate-ncp (ip | ipv6 | dual-stack-passive)
            ipcp-suggest-dns-option;
            mru size;
            mtu (size | use-lower-layer);
            on-demand-ip-address;
            pap;
            peer-ip-address-optional;
        }
        family inet {
            unnumbered-address interface-name;
            address address;
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                output {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
            }
        }
        filter {
            input filter-name {
                precedence precedence;
            }
            output filter-name {
                precedence precedence;
            }
        }
    }
    filter {
        input filter-name;
        output filter-name;
    }

```




```
}
}
```

Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	In a dynamic profile, configure a logical unit number for the dynamic PPPoE logical interface. You must configure a logical interface to be able to use the router.
Options	<p><i>logical-unit-number</i>—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the unit <i>logical-unit-number</i> statement for dynamic PPPoE logical interfaces, you must use the predefined variable <i>\$junos-interface-unit</i> in place of <i>logical-unit-number</i>. The <i>\$junos-interface-unit</i> predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a PPPoE Dynamic Profile • Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

untagged

Syntax	untagged;
Hierarchy Level	[edit interfaces ps0]
Release Information	Statement introduced in Junos OS Release 13.1.
Description	Specify that the router supports untagged traffic on pseudowire subscriber interfaces.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Pseudowire Subscriber Logical Interface Device on page 315

user-group-profile

Syntax	<code>user-group-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit access profile <i>profile-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a configured PPP group profile to PPP users.
<div> NOTE: If <code>user-group-profile</code> is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.</div>	
Options	<i>profile-name</i> —Name of a PPP group profile configured at the <code>[edit access group-profile <i>profile-name</i>]</code> hierarchy level.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying a Configured PPP Group Profile to a Tunnel• Configuring an L2TP Access Profile on the LNS on page 246

version (BFD)

Syntax	version (0 1 automatic);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit system services dhcp-local-server liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server dhcpv6 liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the BFD protocol version to detect.
Options	<p>0—Use BFD protocol version 0.</p> <p>1—Use BFD protocol version 1.</p> <p>automatic—Autodetect the BFD protocol version.</p> <p>Default: automatic</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients on page 94 • Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients on page 89 • Configuring BFD for LDP LSPs


weighted-load-balancing (L2TP LAC)

Syntax	weighted-load-balancing;
Hierarchy Level	[edit services l2tp]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>Specify that the router considers tunnel weight when selecting from among multiple tunnels that share the same preference level. A higher maximum session limit on a tunnel corresponds to a higher tunnel weight. A tunnel with a higher weight is more likely to be selected than a tunnel with a lower weight. The distribution of sessions across all tunnels in the preference level, on average, is proportional to the tunnel weight</p> <p>Disabled by default. By default, tunnel selection within a preference level is strictly random. The destination-equal-load-balancing statement must be disabled to successfully enable this statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 211• Configuring the L2TP LAC Tunnel Selection Parameters on page 209


vlan-id (Dynamic Profiles)

Syntax	<code>vlan-id (<i>number</i> none);</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5. VLAN demux interface support introduced in Junos OS Release 10.2.
Description	For VLAN demux, Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.
Options	<p><i>number</i>—A valid VLAN identifier. When used in the dynamic-profiles hierarchy, specify the \$junos-vlan-id predefined variable to dynamically obtain the VLAN identifier.</p> <p><i>none</i>—Enable the use of untagged pseudo-wire frames on dynamic interfaces.</p> <ul style="list-style-type: none"> For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023. For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles</i>



vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 13.2 for PTX Series Routers. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.
Description	For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
<div>  <p>NOTE: On EX Series switches except for EX4300 and EX9200 switches, the vlan-tagging and family ethernet-switching statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to family ethernet-switching by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default family setting.</p> </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>802.1Q VLANs Overview</i> • <i>Configuring a Layer 3 Subinterface (CLI Procedure)</i> • <i>Configuring Tagged Aggregated Ethernet Interfaces</i> • <i>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i>

vlan-tagging (Dynamic)

Syntax	vlan-tagging;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
<div>  <p>NOTE: For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.</p> </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs</i> • <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i> • Configuring the L2TP LNS Peer Interface on page 251

vlan-tags

Syntax	<code>vlan-tags outer [<i>tpid</i>].<i>vlan-id</i> [inner [<i>tpid</i>].<i>vlan-id</i>];</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5. VLAN demux interface support introduced in Junos OS Release 10.2.
Description	For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the stacked-vlan-tagging statement at the <code>[edit interfaces <i>interface-name</i>]</code> hierarchy level.
<div>  NOTE: The inner-range <i>vid1–vid2</i> option is supported on IQE PICs only. </div>	
Options	inner [<i>tpid</i>].<i>vlan-id</i> —A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i> . When used in the dynamic-profiles hierarchy, specify the <code>\$junos-vlan-id</code> predefined variable to dynamically obtain the VLAN ID.
<div>  NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the inner-range <i>tpid. vid1–vid2</i> option with the vlan-tags statement for ISP-facing interfaces. </div>	
Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.	
outer [<i>tpid</i>].<i>vlan-id</i> —A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i> . When used in the dynamic-profiles hierarchy, specify the <code>\$junos-stacked-vlan-id</code> predefined variable.	
Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring Dual VLAN Tags

CHAPTER 38

Operational Commands

- clear services l2tp destination
- clear services l2tp destination lockout
- clear services l2tp session
- clear services l2tp session statistics
- clear services l2tp tunnel
- clear services l2tp tunnel statistics
- request interface (revert | switchover) (Aggregated Inline Service Interfaces)
- restart
- show bfd subscriber session
- show dynamic-profile session
- show interfaces ps0 (Pseudowire Subscriber Interfaces)
- show interfaces redundancy
- show ppp interface
- show ppp statistics
- show ppp summary
- show services inline ip-reassembly statistics
- show services l2tp client
- show services l2tp destination
- show services l2tp destination lockout
- show services l2tp session
- show services l2tp session-limit-group
- show services l2tp summary
- show services l2tp tunnel
- show services l2tp tunnel-group
- show services l2tp tunnel-switch destination
- show services l2tp tunnel-switch session
- show services l2tp tunnel-switch summary
- show services l2tp tunnel-switch tunnel

- `show services soft-gre tunnel`
- `show subscribers`
- `show subscribers summary`
- `show system subscriber-management statistics`
- `show system subscriber-management summary`
- `test services l2tp tunnel`

clear services l2tp destination

Syntax clear services l2tp destination
<all | local-gateway *gateway-address* | peer-gateway *gateway-address*>

Release Information Command introduced in Junos OS Release 10.4.

Description Clear all Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers.



NOTE: You cannot issue the clear services l2tp destination command in parallel with statistics-related show services l2tp commands from separate terminals. If this clear command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the show commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics
show services l2tp destination statistics	show services l2tp tunnel extensive
show services l2tp session extensive	show services l2tp tunnel statistics
show services l2tp session statistics	

Options all—Close all L2TP destinations.



BEST PRACTICE: The all option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the all option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

local-gateway *gateway-address*—Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.

peer-gateway *gateway-address*—Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.

Required Privilege Level clear

Related Documentation • [show services l2tp destination on page 732](#)

List of Sample Output [clear services l2tp destination all on page 664](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear services l2tp destination all](#)

```
user@host> clear services l2tp destination all  
  
Destination 2 closed
```

clear services l2tp destination lockdown

Syntax clear services l2tp destination lockdown
<all | local-gateway *gateway-address* | peer-gateway *gateway-address*>

Release Information Command introduced in Junos OS Release 16.1 on MX Series Routers.

Description Clear the lockdown timer for all or only the specified Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. Clearing the lockdown timer removes the destination from the lockdown list. This command is available only for LAC on MX Series routers.



NOTE: You cannot issue the clear services l2tp destination command in parallel with statistics-related show services l2tp commands from separate terminals. If this clear command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the show commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics
show services l2tp destination statistics	show services l2tp tunnel extensive
show services l2tp session extensive	show services l2tp tunnel statistics
show services l2tp session statistics	

Options all—(Optional) Unlock all L2TP destinations.

local-gateway *gateway-address*—(Optional) Unlock only the L2TP destination with the specified local gateway address.

peer-gateway *gateway-address*—(Optional) Unlock only the L2TP destination with the specified address.

Required Privilege Level clear

Related Documentation

- [clear services l2tp destination on page 663](#)
- [show services l2tp destination on page 732](#)

List of Sample Output [clear services l2tp destination lockdown all on page 666](#)

Output Fields When you enter this command, you are provided no feedback on the status of your request.

Sample Output

`clear services l2tp destination lockout all`

```
user@host> clear services l2tp destination lockout all
```

clear services l2tp session

Syntax clear services l2tp session (all | 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* | tunnel-group *group-name* | user *username*)

Release Information Command introduced before Junos OS Release 7.4.

Description (M10i and M7i routers only) Clear Layer 2 Tunneling Protocol (L2TP) sessions on LNS.
(MX Series routers only) Clear L2TP sessions on LAC and LNS.



NOTE: On MX Series routers, you cannot issue the clear services l2tp session command in parallel with statistics-related show services l2tp commands from separate terminals. If this clear command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the show commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics
show services l2tp destination statistics	show services l2tp tunnel extensive
show services l2tp session extensive	show services l2tp tunnel statistics
show services l2tp session statistics	

Options all—Close all L2TP sessions.



BEST PRACTICE: The all option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the all option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

interface *interface-name*—Clear only the L2TP sessions using 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—Clear only the L2TP sessions associated with the specified local gateway address.

local-gateway-name gateway-name—Clear only the L2TP sessions associated with the specified local gateway name.

local-session-id session-id—Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id tunnel-id—Clear only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway gateway-address—Clear only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name gateway-name—Clear only the L2TP sessions associated with the peer gateway with the specified name.

tunnel-group group-name—Clear only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

user username—(M Series routers only) Clear only the L2TP sessions for the specified username.

Required Privilege Level clear

Related Documentation

- [L2TP Services Configuration Overview](#)
- [L2TP Minimum Configuration](#)
- [clear services l2tp session statistics on page 670](#)
- [show services l2tp session on page 738](#)

List of Sample Output [clear services l2tp session on page 668](#)
[clear services l2tp session interface on page 669](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear services l2tp session](#)

```
user@host> clear services l2tp session 31694
```

```
Session 31694 closed
```


Sample Output

clear services l2tp session interface

```
user@host> show services l2tp session Tunnel local ID: 17185
```

Local ID	Remote ID	State	Interface unit	Interface Name
5117	1	Established	1073741828	si-2/0/0
34915	2	Established	1073741829	si-2/1/0
6454	3	Established	1073741830	si-2/0/0
46142	4	Established	1073741831	si-2/1/0

```
user@host> clear services l2tp session interface si-2/0/0
```

```
Session 5117 closed
Session 6454 closed
```

```
user@host> show services l2tp session Tunnel local ID: 17185
```

Local ID	Remote ID	State	Interface unit	Interface Name
34915	2	Established	1073741829	si-2/1/0
46142	4	Established	1073741831	si-2/1/0

clear services l2tp session statistics

Syntax	<code>clear services l2tp session statistics (all interface <i>interface-name</i> local-gateway <i>gateway-address</i> local-gateway-name <i>gateway-name</i> local-session-id <i>session-id</i> local-tunnel-id <i>tunnel-id</i> peer-gateway <i>gateway-address</i> peer-gateway-name <i>gateway-name</i> tunnel-group <i>group-name</i> user <i>username</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Support for MX Series routers added in Junos OS Release 10.4.
Description	(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions.
Options	<p>all—Clear statistics for all L2TP sessions.</p> <p>interface <i>interface-name</i>—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:</p> <ul style="list-style-type: none">• si-<i>fpc/pic/port</i>—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.• sp-<i>fpc/pic/port</i>—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers. <p>local-gateway <i>gateway-address</i>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.</p> <p>local-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.</p> <p>local-session-id <i>session-id</i>—Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.</p> <p>local-tunnel-id <i>tunnel-id</i>—Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.</p> <p>peer-gateway <i>gateway-address</i>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.</p> <p>peer-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.</p> <p>tunnel-group <i>group-name</i>—Clear statistics for only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p> <p>user <i>username</i>—Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.</p>

Required Privilege Level view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- [clear services l2tp session on page 667](#)
- [show services l2tp session on page 738](#)

List of Sample Output [clear services l2tp session statistics all on page 671](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear services l2tp session statistics all](#)

```
user@host> clear services l2tp session statistics all
Session 26497 statistics cleared
```

clear services l2tp tunnel

Syntax clear services l2tp tunnel (all | interface *sp-fpc/pic/port* | local-gateway *gateway-address* | local-gateway-name *gateway-name* | local-tunnel-id *tunnel-id* | peer-gateway *gateway-address* | peer-gateway-name *gateway-name* | tunnel-group *group-name*)

Release Information Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description (M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear Layer 2 Tunneling Protocol (L2TP) tunnels.



NOTE: On MX Series routers, you cannot issue the clear services l2tp tunnel command in parallel with statistics-related show services l2tp commands from separate terminals. If this clear command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the show commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics
show services l2tp destination statistics	show services l2tp tunnel extensive
show services l2tp session extensive	show services l2tp tunnel statistics
show services l2tp session statistics	

Options all—Clear all L2TP tunnels.



BEST PRACTICE: The all option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the all option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

sp-fpc/pic/port—(Optional) Clear only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP on MX Series routers.

local-gateway gateway-address—Clear only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*—Clear only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*—Clear only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear only the L2TP tunnels associated with the peer gateway with the specified name.

tunnel-group *group-name*—Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- [clear services l2tp tunnel statistics on page 674](#)
- [show services l2tp tunnel on page 755](#)

List of Sample Output [clear services l2tp tunnel on page 673](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear services l2tp tunnel`

```
user@host> clear services l2tp tunnel 17185

Tunnel 17185 closed
```

clear services l2tp tunnel statistics

Syntax	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>)
Release Information	Command introduced before Junos OS Release 7.4. Support for MX Series routers added in Junos OS Release 10.4.
Description	(M10i and M7i routers: LNS only. MX Series routers: LAC only.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels.
Options	<p>all—Clear statistics for all L2TP tunnels.</p> <p>interface <i>sp-fpc/pic/port</i>—Clear statistics for only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP LAC on MX Series routers.</p> <p>local-gateway <i>gateway-address</i>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.</p> <p>local-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.</p> <p>local-tunnel-id <i>tunnel-id</i>—Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.</p> <p>peer-gateway <i>gateway-address</i>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.</p> <p>peer-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.</p> <p>tunnel-group <i>group-name</i>—Clear statistics for only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• L2TP Services Configuration Overview• L2TP Minimum Configuration• clear services l2tp tunnel on page 672• show services l2tp tunnel on page 755
List of Sample Output	clear services l2tp tunnel statistics all on page 675

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear services l2tp tunnel statistics all`

```
user@host> clear services l2tp tunnel statistics all
Tunnel 9933 statistics cleared
```

request interface (revert | switchover) (Aggregated Inline Service Interfaces)

Syntax request interface (revert | switchover) *bundle-name*

Release Information Command introduced in Junos OS Release 16.2.

Description Manually revert L2TP data traffic from the designated backup link to the designated primary link of an aggregated inline service interface bundle interface for which 1:1 redundancy is configured, or manually switch data traffic from the primary link to the backup link.



NOTE: When 1:1 redundancy protection is configured for an aggregated inline service interface, if the primary link fails, the router automatically routes data traffic destined for the L2TP session on that link to the backup link. However, the router does not automatically route data traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually divert traffic back to the primary link by issuing the `request interface revert` operational command.

Options **revert**—Restore data traffic for the LNS session to the primary link.

switchover—Transfer data traffic for the LNS session to the secondary (backup) link.

bundle-name—Name of the aggregated inline service interface bundle.

Required Privilege Level view

List of Sample Output [request interface switchover on page 676](#)
[request interface revert on page 676](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`request interface switchover`

```
user@host >request interface switchover asi0
error: requesting cmd SWITCH when primary is not active
```

Sample Output

`request interface revert`

```
user@host >request interface revert asi0
request succeeded
```


restart

- List of Syntax**
- Syntax on page 678
 - Syntax (ACX Series Routers) on page 678
 - Syntax (EX Series Switches) on page 678
 - Syntax (MX Series Routers) on page 679
 - Syntax (QFX Series) on page 679
 - Syntax (Routing Matrix) on page 679
 - Syntax (TX Matrix Routers) on page 679
 - Syntax (TX Matrix Plus Routers) on page 680
 - Syntax (QFX Series) on page 680

Syntax restart

```
<adaptive-services | ancpd-service | application-identification | audit-process |
auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
class-of-service | clksyncd-service | database-replication | datapath-trace-service
| dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
ecc-error-logging | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | iccp-service | idp-policy | immediately
| interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
| l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
| local-policy-decision-function | mac-validation | mib-process | mounstd-service
| mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
packet-triggered-subscribers | peer-selection-service | pgm | pic-services-logging | pki-service
| ppp | ppp-service | pppoe | protected-system-domain-service |
redundancy-interface-process | remote-operations | root-system-domain-service | routing
<logical-system logical-system-name> | sampling | sbc-configuration-process | sdk-service
| service-deployment | services | snmp | soft | static-subscribers | statistics-service |
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
vrrp | web-management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers) restart

```
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
| disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | immediately | interface-control |
ipsec-key-management | l2-learning | lacp | link-management | mib-process | mounstd-service
| mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service | ppp | pppoe |
redundancy-interface-process | remote-operations | routing | sampling | sdk-service
| secure-neighbor-discovery | service-deployment | services | snmp | soft | statistics-service |
subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>
```

Syntax (EX Series Switches) restart

```
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
ethernet-switching | event-processing | firewall | general-authentication-service |
interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
| lldpd-service | mib-process | mounstd-service | multicast-snooping | pgm |
```

redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery
| service-deployment | sflow-service | snmp | vrrp | web-management>

Syntax (MX Series Routers) restart
 <adaptive-services | ancpd-service | application-identification | audit-process |
 auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
 class-of-service | clksyncd-service | database-replication | datapath-trace-service
 | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
 ecc-error-logging | ethernet-connectivity-fault-management
 | ethernet-link-fault-management | event-processing | firewall |
 general-authentication-service | gracefully | iccp-service | idp-policy | immediately
 | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
 | l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
 | local-policy-decision-function | mac-validation | mib-process | mounstd-service
 | mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
 packet-triggered-subscribers | peer-selection-service | pgm | pic-services-logging |
 pki-service | ppp | ppp-service | pppoe | protected-system-domain-service |
 redundancy-interface-process | remote-operations | root-system-domain-service | routing
 | routing <logical-system *logical-system-name*> | sampling | sbc-configuration-process |
 sdk-service | service-deployment | services | snmp | soft | static-subscribers | statistics-service |
 subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
 vrrp | web-management>
 <all-members>
 <gracefully | immediately | soft>
 <local>
 <member *member-id*>

Syntax (QFX Series) restart
 <adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
 diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall
 | general-authentication-service | igmp-host-services | interface-control |
 ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
 named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
 redundancy-interface-process | remote-operations | *logical-system-name*> | routing |
 sampling | secure-neighbor-discovery | service-deployment | snmp | usb-control |
 web-management>
 <gracefully | immediately | soft>

Syntax (Routing Matrix) restart
 <adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring |
 dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control
 | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp |
 link-management | mib-process | pgm | pic-services-logging | ppp | pppoe |
 redundancy-interface-process | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp>
 <all | all-lcc | lcc *number*>
 <gracefully | immediately | soft>

Syntax (TX Matrix Routers) restart
 <adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service |
 diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging |
 event-processing | firewall | interface-control | ipsec-key-management | kernel-replication
 | l2-learning | l2tp-service | lacp | link-management | mib-process | pgm | pic-services-logging

	<p> ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p> <p><all-chassis all-lcc lcc <i>number</i> scc></p> <p><gracefully immediately soft></p>
Syntax (TX Matrix Plus Routers)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p> <p><all-chassis all-lcc all-sfc lcc <i>number</i> sfc <i>number</i>></p> <p><gracefully immediately soft></p>
Syntax (QFX Series)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dialer-services diameter-service dlsf ethernet-connectivity event-processing fibre-channel firewall general-authentication-service igmp-host-services interface-control ipsec-key-management isdn-signaling l2ald l2-learning l2tp-service mib-process named-service network-access-service nstrace-process pgm ppp pppoe redundancy-interface-process remote-operations <i>logical-system-name</i>> routing sampling secure-neighbor-discovery service-deployment snmp usb-control web-management></p> <p><gracefully immediately soft></p>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Options added:</p> <ul style="list-style-type: none"> • dynamic-flow-capture in Junos OS Release 7.4. • dlsf in Junos OS Release 7.5. • event-processing in Junos OS Release 7.5. • ppp in Junos OS Release 7.5. • l2ald in Junos OS Release 8.0. • link-management in Release 8.0. • pgcp-service in Junos OS Release 8.4. • sbc-configuration-process in Junos OS Release 9.5. • services pgcp gateway in Junos OS Release 9.6. • sfc and all-sfc for the TX Matrix Router in Junos OS Release 9.6.
Description	Restart a Junos OS process.



CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

dapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dls—(QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lACP—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

moundd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(MX Series routers) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation • *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output [restart interface-control gracefully on page 687](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[restart interface-control gracefully](#)

```
user@host> restart interface-control gracefully
Interface control process started, pid 41129
```

show bfd subscriber session

Syntax	show bfd subscriber session <brief detail extensive summary>
Release Information	Command introduced in Junos OS Release 15.1 on MX Series routers.
Description	Display information about active Bidirectional Forwarding Detection (BFD) subscriber sessions.
Options	none —(Same as brief) Display information about active BFD subscriber sessions. brief detail extensive summary —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>clear bfd session</i> • <i>Examples: Configuring BFD for Static Routes</i>
List of Sample Output	show bfd subscriber session on page 690 show bfd subscriber session brief on page 690 show bfd subscriber session detail on page 691 show bfd subscriber session extensive on page 691 show bfd subscriber session summary on page 691
Output Fields	Table 28 on page 688 describes the output fields for the show bfd subscriber session command. Output fields are listed in the approximate order in which they appear.

Table 28: show bfd subscriber session Output Fields

Field Name	Field Description	Level of Output
Address	IP Address on which the BFD subscriber session is active.	brief detail extensive none
State	State of the BFD subscriber session: Up , Down , Init (initializing), or Failing .	brief detail extensive none
Interface	Interface on which the BFD subscriber session is active.	brief detail extensive none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	brief detail extensive none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	brief detail extensive none

Table 28: show bfd subscriber session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	detail extensive
Session up time	How long a BFD subscriber session has been established in <i>hh:mm:ss</i> .	detail extensive
Client	Protocol for which the BFD subscriber session is active: DHCP , ISIS , OSPF , or Static .	detail extensive
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	detail extensive
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	detail extensive
Local diagnostic	Local diagnostic information about failing BFD subscriber sessions.	detail extensive
Remote diagnostic	Remote diagnostic information about failing BFD subscriber sessions.	detail extensive
Remote state	Indication that the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	detail extensive
Version	BFD version: 0 or 1 .	extensive
Replicated	Indication that nonstop routing or graceful Routing Engine switchover is configured and the BFD subscriber session has been replicated to the backup Routing Engine.	detail extensive
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD subscriber session.	extensive
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD subscriber session.	extensive
Adaptive async TX interval	Transmission interval being used because of adaptation.	extensive
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	extensive
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	extensive
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	extensive
Remote min RX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	extensive

Table 28: show bfd subscriber session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local discriminator	Authentication code used by the local system to identify that BFD subscriber session.	extensive
Remote discriminator	Authentication code used by the remote system to identify that BFD subscriber session.	extensive
Echo mode	Information about the state of echo transmissions on the BFD subscriber session, such as disabled or inactive.	extensive
Remote is control-plane independent	Indication that the BFD subscriber session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart. The following BFD subscriber sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.	extensive
Session ID	BFD subscriber session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	detail extensive
sessions	Total number of active BFD subscriber sessions.	All levels
clients	Total number of clients that are hosting active BFD subscriber sessions.	All levels
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	detail extensive
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	detail extensive

Sample Output

show bfd subscriber session

```

user@host> show bfd subscriber session

              Detect      Transmit
Address      State      Interface  Time      Interval  Multiplier
203.0.113.2   Up        ae0.0      90.000    30.000     3
203.0.113.6   Up        ae0.1      90.00    30.000     3
203.0.113.10  Up        ae0.2      90.000    30.000     3
203.0.113.14  Up        ae0.3      90.000    30.000     3
203.0.113.18  Up        ae0.4      90.000    30.000     3

20 sessions, 20 clients

```

show bfd subscriber session brief

The output for the **show bfd subscriber session brief** command is identical to that for the **show bfd subscriber session** command.

show bfd subscriber session detail

```

user@host> show bfd subscriber session detail

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
203.0.113.2	Up	ae0.0	90.000	30.000	3
Client DHCP, TX interval 30.000, RX interval 30.000, multiplier 3					
Session up time 09:11:50					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
203.0.113.6	Up	ae0.1	90.000	30.000	3
Client DHCP, TX interval 30.000, RX interval 30.000					
Session up time 09:11:50					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					

20 sessions, 20 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd subscriber session extensive

```

user@host> show bfd subscriber session extensive

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
203.0.113.2	Up	ae0.0	90.000	30.000	3
Client DHCP, TX interval 30.000, RX interval 30.000					
Session up time 09:11:50					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
Min async interval 30.000, min slow interval 30.000					
Adaptive async TX interval 30.000, RX interval 30.000					
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3					
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3					
Local discriminator 20, remote discriminator 16					
Echo mode disabled/inactive					
Remote is control-plane independent					
Session ID: 0x1					
203.0.113.6	Up	ae0.1	90.000	30.000	3
Client DHCP, TX interval 30.000, RX interval 30.000					
Session up time 09:11:50					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
Min async interval 30.000, min slow interval 30.000					
Adaptive async TX interval 30.000, RX interval 30.000					
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3					
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3					
Local discriminator 21, remote discriminator 17					
Echo mode disabled/inactive					
Remote is control-plane independent					
Session ID: 0x2					

show bfd subscriber session summary

```

user@host> show bfd subscriber session summary

```

20 sessions, 20 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show dynamic-profile session

Syntax show dynamic-profile session
 <client-id *client-id*>
 <profile-name *profile-name*>
 <service-id *service-id*>

Release Information Command introduced in Junos OS Release 13.3.

Description Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.



NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as **NONE** in the command output.
- The variable is also displayed as **NONE** when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the **show configuration dynamic-profiles** command.

Options **client-id *client-id***—Display dynamic profile information for subscribers associated with the specified client.

profile-name *profile-name*—(Optional) Display dynamic profile information for the specified subscriber or service profile.

service-id *service-id*—Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level view

List of Sample Output [show dynamic-profile session client-id \(Client ID\) on page 694](#)
[show dynamic-profile session client-id profile-name \(Client ID and Dynamic Profile\) on page 696](#)

[show dynamic-profile session service-id \(Service Session\) on page 696](#)

Output Fields This command displays the dynamic client or service profile configuration for each subscriber.

Sample Output

[show dynamic-profile session client-id \(Client ID\)](#)

```
user@host>show dynamic-profile session client-id 20
pppoe {
  interfaces {
    pp0 {
      unit 1073741831 {
        ppp-options {
          chap;
          pap;
        }
        pppoe-options {
          underlying-interface ge-2/0/0.0;
          server;
        }
        family {
          inet {
            unnumbered-address lo0.0;
          }
        }
      }
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tcp1 {
      scheduler-map smap1_UID1024;
      shaping-rate 100m;
    }
  }
  interfaces {
    pp0 {
      unit 1073741831 {
        output-traffic-control-profile tcp1;
      }
    }
  }
  scheduler-maps {
    smap1_UID1024 {
      forwarding-class best-effort scheduler sch1_UID1023;
    }
  }
  schedulers {
    sch1_UID1023 {
      transmit-rate percent 40;
      buffer-size percent 40;
      priority low;
    }
  }
}
}
filter-service {
```

```

interfaces {
  pp0 {
    unit 1073741831 {
      family {
        inet {
          filter {
            input input-filter_UID1026 precedence 50;
            output output-filter_UID1027 precedence 50;
          }
        }
      }
    }
  }
}
firewall {
  family {
    inet {
      filter input-filter_UID1026 {
        interface-specific;
        term t1 {
          then {
            policer policer1_UID1025;
            service-accounting;
          }
        }
        term rest {
          then accept;
        }
      }
      filter output-filter_UID1027 {
        interface-specific;
        term rest {
          then accept;
        }
      }
    }
  }
  policer policer1_UID1025 {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 15k;
    }
    then discard;
  }
}
}
cos-service {
  class-of-service {
    scheduler-maps {
      smap2_UID1029 {
        forwarding-class assured-forwarding scheduler sch2_UID1028;
      }
    }
    schedulers {
      sch2_UID1028 {
        transmit-rate percent 60;
        buffer-size percent 60;
        priority high;
      }
    }
  }
}
}

```

```
}  
bsimmons  
}
```

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile)

```
user@host>show dynamic-profile session client-id 20 profile-name cos-service  
cos-service {  
  class-of-service {  
    scheduler-maps {  
      smap2_UID1029 {  
        forwarding-class assured-forwarding scheduler sch2_UID1028;  
      }  
    }  
    schedulers {  
      sch2_UID1028 {  
        transmit-rate percent 60;  
        buffer-size percent 60;  
        priority high;  
      }  
    }  
  }  
}
```

show dynamic-profile session service-id (Service Session)

```
user@host>show dynamic-profile session service-id 21  
filter-service {  
  interfaces {  
    pp0 {  
      unit 1073741831 {  
        family {  
          inet {  
            filter {  
              input input-filter_UID1026 precedence 50;  
              output output-filter_UID1027 precedence 50;  
            }  
          }  
        }  
      }  
    }  
  }  
}  
firewall {  
  family {  
    inet {  
      filter input-filter_UID1026 {  
        interface-specific;  
        term t1 {  
          then {  
            policer policer1_UID1025;  
            service-accounting;  
          }  
        }  
        term rest {  
          then accept;  
        }  
      }  
      filter output-filter_UID1027 {  
        interface-specific;  
      }  
    }  
  }  
}
```

```
        term rest {
            then accept;
        }
    }
}
policer policer1_UID1025 {
    if-exceeding {
        bandwidth-limit 1m;
        burst-size-limit 15k;
    }
    then discard;
}
}
```

show interfaces ps0 (Pseudowire Subscriber Interfaces)

Syntax	show interfaces ps0 <brief detail extensive terse>
Release Information	Command introduced at Junos OS Release 13.1.
Description	Display status information about the pseudowire subscriber interface.
Options	brief detail extensive terse —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Pseudowire Subscriber Logical Interfaces Overview on page 307
List of Sample Output	show interfaces ps0 on page 700 show interfaces ps0 extensive on page 701
Output Fields	Table 29 on page 698 lists the output fields for the show interfaces ps0 command. Output fields are listed in the approximate order in which they appear.

Table 29: show interfaces ps0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	brief detail extensive none
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	brief detail extensive none
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Physical interface type (Software-Pseudo).	brief detail extensive none
Link-level type	Encapsulation being used on the physical interface.	brief detail extensive
MTU	MTU size on the physical interface.	brief detail extensive
Clocking	Reference clock source. It can be Internal or External .	brief detail extensive
Speed	Speed at which the interface is running.	brief detail extensive

Table 29: show interfaces ps0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	brief detail extensive none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	brief detail extensive none
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up or up to down. The format is Last flapped: <i>year-month-day hours:minutes:seconds: timezone (hours:minutes:seconds ago)</i> . or Never. For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).	detail extensive none
input packets	Number of packets received on the logical interface.	detail extensive none
output packets	Number of packets transmitted on the logical interface.	detail extensive none
Logical Interface		
Logical interface	Name of the logical interface.	brief detail extensive none
Index	Logical interface index number (which reflects its initialization sequence).	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	brief detail extensive none
Encapsulation	Type of encapsulation configured on the logical interface.	brief extensive none
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	detail extensive

Table 29: show interfaces ps0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	detail extensive
Transit statistics	<p>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. This counter usually takes less than 1 second to stabilize.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>	detail extensive
Protocol	Protocol family configured on the logical interface.	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Addresses, Flags	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive terse none
Broadcast	Broadcast address.	detail extensive none

Sample Output

show interfaces ps0

```
user@host> show interfaces ps0
```



```

Physical interface: ps0, Enabled, Physical link is Up
Interface index: 166, SNMP ifIndex: 658
Type: Software-Pseudo, Link-level type: 90, MTU: 1518, Clocking: 1, Speed: 800mbps

Device flags : Present Running
Interface flags: Point-To-Point Internal: 0x4000
Current address: 00:00:5E:00:53:4a, Hardware address: 00:00:5E:00:53:4a
Last flapped : Never
Input packets : 0
Output packets: 0

Logical interface ps0.0 (Index 74) (SNMP ifIndex 656)
Flags: Point-To-Point 0x4000 Encapsulation: Ethernet-CCC
Input packets : 482
Output packets: 0
Protocol ccc, MTU: 1518
Flags: Is-Primary

Logical interface ps0.1 (Index 78) (SNMP ifIndex 665)
Flags: Point-To-Point 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Input packets : 0
Output packets: 482
Protocol inet, MTU: 1500
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 203.0.113.0/24, Local: 203.0.113.1, Broadcast: 203.0.113.255

Logical interface ps0.32767 (Index 75) (SNMP ifIndex 692)
Flags: Point-To-Point 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

show interfaces ps0 extensive

```

user@host> show interfaces ps0.1 extensive
Logical interface ps0.1 (Index 389) (SNMP ifIndex 0) (Generation 199)
Flags: Up 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
IPv6 transit 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
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0

```

```
Output packets: 0
Protocol inet, MTU: 1500, Generation: 194, Route table: 0
  Flags: Sendbcst-pkt-to-re, Unnumbered
  Donor interface: lo0.0 (Index 322)
  Addresses, Flags: Primary Is-Default Is-Primary
    Destination: Unspecified, Local: 203.0.113.144, Broadcast: Unspecified,
Generation: 138
Protocol inet6, MTU: 1500, Generation: 198, Route table: 0
  Flags: Unnumbered
  Donor interface: lo0.0 (Index 322)
    Destination: Unspecified, Local: 2001:db8::e187
Generation: 157
```

show interfaces redundancy

Syntax `show interfaces redundancy`
`<brief | detail>`

Release Information Command introduced before Junos OS Release 7.4.
detail option added in Junos OS Release 10.0.

Description (M Series, T Series, and MX Series routers only) Display general information about redundancy for aggregated multiservices (AMS) interfaces configured for warm standby, adaptive services and link services intelligent queuing (IQ) interfaces, aggregated Ethernet interfaces redundancy, and LNS aggregated inline service interfaces.



NOTE: When you run the `show interfaces redundancy` command on an MX80 router, it displays the error message, `error:the redundancy-interface-process subsystem is not running`. This is because an MX80 router does not have a redundant FPC and does not support link protection.

Options `brief | detail`—(Optional) Display the specified level of output.

Required Privilege Level view

List of Sample Output [show interfaces redundancy on page 704](#)
[show interfaces redundancy \(Aggregated Ethernet\) on page 704](#)
[show interfaces redundancy \(Aggregated Inline Service Interface\) on page 704](#)
[show interfaces redundancy detail on page 704](#)

Output Fields [Table 30 on page 703](#) lists the output fields for the `show interfaces redundancy` command. Output fields are listed in the approximate order in which they appear.

Table 30: show interfaces redundancy Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the AMS interface, redundant adaptive services, link services IQ interfaces, aggregated Ethernet interfaces, or LNS aggregated inline service interfaces.	All levels
State	State of the redundant interface: Not present , On primary , On secondary , or Waiting for primary MS PIC .	All levels

Table 30: show interfaces redundancy Output Fields (*continued*)

Field Name	Field Description	Level of Output
Last Change	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: <ul style="list-style-type: none"> • GRES is not configured on the router. • The rlsq interface is configured without the hot-standby or warm-standby statements and the backup lsq interface was active before the switchover. • No logical interfaces are configured or all of the configured logical interfaces are down at the time of the switchover. 	All levels
Primary	Name of the interface configured to be the primary interface.	All levels
Secondary	Name of the interface configured to be the backup interface.	All levels
Current Status	Physical status of the primary and secondary interfaces.	All levels
Mode	Standby mode.	detail

Sample Output

show interfaces redundancy

```

user@host> show interfaces redundancy
Interface State      Last change Primary Secondary Current status
rsp0      Not present
rsp1      On secondary 1d 23:56 sp-1/2/0 sp-0/3/0 primary down
rsp2      On primary 10:10:27 sp-1/3/0 sp-0/2/0 secondary down
rlsq0     On primary 00:06:24 lsq-0/3/0 lsq-1/0/0 both up
ams0      On primary 00:39:51 mams-5/0/0 mams-5/1/0 both up

```

show interfaces redundancy (Aggregated Ethernet)

```

user@host> show interfaces redundancy
Interface State      Last change Primary Secondary Current status
rlsq0     On secondary 00:56:12 lsq-4/0/0 lsq-3/0/0 both up

ae0
ae1
ae2
ae3
ae4

```

show interfaces redundancy (Aggregated Inline Service Interface)

```

user@host> show interfaces redundancy asi0
Interface State      Last change Primary Secondary Current status
asi0      On primary 00:00:09 si-1/0/0 si-0/0/0 both up

```

show interfaces redundancy detail

```

user@host> show interfaces redundancy detail

```

Interface : rlsq0
State : On primary
Last change : 00:45:47
Primary : lsq-0/2/0
Secondary : lsq-1/2/0
Current status : both up
Mode : hot-standby

Interface : rlsq0:0
State : On primary
Last change : 00:45:46
Primary : lsq-0/2/0:0
Secondary : lsq-1/2/0:0
Current status : both up
Mode : warm-standby

Interface : asi0
State : On primary
Last change : 00:03:42
Primary : si-1/0/0
Secondary : si-0/0/0
Mode : hot-standby
Current status : both up

Interface :ams0
State : On primary
Last change : 00:39:52
Primary : mams-5/0/0
Secondary : mams-5/1/0
Mode : warm-standby
Current status : both up
Replication state : Disconnected

show ppp interface

Syntax	<code>show ppp interface <i>interface-name</i></code> <code><extensive terse></code>
Release Information	Command introduced in Junos OS Release 7.5.
Description	Display information about PPP interfaces.
Options	<i>interface-name</i> —Name of a logical interface. extensive terse —(Optional) Display the specified level of output.
Required Privilege Level	view
List of Sample Output	show ppp interface on page 714 show ppp interface extensive on page 714 show ppp interface terse on page 714
Output Fields	Table 31 on page 706 lists the output fields for the show ppp interface command. Output fields are listed in the approximate order in which they appear.

Table 31: show ppp interface Output Fields

Field Name	Field Description	Level of Output
Session	Name of the logical interface on which the session is running.	All levels
Type	Session type: PPP.	All levels
Phase	PPP process phase: Authenticate , Pending , Establish , LCP , Network , Disabled , and Tunneled .	All levels
Session flags	Special conditions present in the session: Bundled , TCC , No-keepalives , Looped , Monitored , and NCP-only .	All levels
protocol State	Protocol state information. See specific protocol state fields for information.	None specified
AUTHENTICATION	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the Authentication field description for further information.	None specified

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive settings	<p>Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS based PPP sessions are supported only on service interfaces (si).</p> <ul style="list-style-type: none"> • Interval—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the interval and Down-count values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • Up-count—The number of keepalive packets a destination must receive to change a link's status from down to up. • Down-count—The number of keepalive packets a destination must fail to receive before the network takes down a link. 	extensive
Magic-Number validation	<p>Indicates whether the local peer is configured to ignore mismatches between peer magic numbers when the numbers are validated during PPP keepalive (Echo-Request/Echo-Reply) exchanges.</p> <ul style="list-style-type: none"> • Enable—Mismatch detection sends failed Echo-Reply packets to the Routing Engine. If a valid magic number is not received within the configurable keepalive interval, PPP treats this as a keepalive failure and tears down the PPP sessions. • Disable—The Packet Forwarding Engine does not perform a validation check for magic numbers received from remote peers. A mismatch cannot be detected, so receipt of its own magic number or an unexpected value does not trigger notification to the Routing Engine. 	extensive
RE Keepalive statistics	<p>Keepalive statistics for the packets handled by the Routing Engine.</p> <ul style="list-style-type: none"> • LCP echo req Tx—LCP echo requests sent from the Routing Engine. • LCP echo req Rx—LCP echo requests received at the Routing Engine. • LCP echo rep Tx—LCP echo responses sent from the Routing Engine. • LCP echo rep Rx—LCP echo responses received at the Routing Engine. • LCP echo req timeout—Number of keepalive packets where the keepalive aging timer has expired. • LCP Rx echo req Magic Num Failures—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • LCP Rx echo rep Magic Num Failures—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match. 	extensive

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
LCP	<p>LCP information:</p> <ul style="list-style-type: none"> • State—LCP protocol state (all platforms except M120 and M320 routers): <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—LCP protocol state (M120 and M320 routers): <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—LCP state start time. • Last completed—LCP state completion time. 	extensive

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • Negotiated options: <ul style="list-style-type: none"> • ACFC—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields. • Asynchronous map—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link. • Authentication protocol—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required. • Authentication algorithm—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported. • Endpoint discriminator class—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link. • Magic number—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated. • MRU—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets. • MRRU—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets. • Multilink header suspendable classes—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given. • Multilink header format classes—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number. • PFC—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field. • short sequence—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers. 	

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Authentication	<p>CHAP or PAP authentication state information. For CHAP authentication:</p> <ul style="list-style-type: none"> • Chap-ans-rcvd—Packet was sent from the peer, indicating that the peer received the Chap-resp-sent packet. • Chap-ans-sent—Packet was sent from the authenticator, indicating that the authenticator received the peer's Chap-resp-rcvd packet. • Chap-chal-rcvd—Challenge packet has been received by the peer. • Chap-chal-sent—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • Chap-resp-rcvd—CHAP response packet has been received by the authenticator. • Chap-resp-sent—CHAP response packet has been sent to the authenticator. • Closed—Link is not available for authentication. • Failure—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • Success—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. <p>For PAP authentication:</p> <ul style="list-style-type: none"> • Pap-resp-sent—PAP response sent to peer (ACK/NACK). • Pap-req-rcvd—PAP request packet received from peer. • Pap-resp-rcvd—PAP response received from the peer (ACK/NACK). • Pap-req-sent—PAP request packet sent to the peer. • Closed—Link is not available for authentication. • Failure—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • Success—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. 	None specified

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPCP	<p>Internet Protocol Control Protocol (IPCP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—IPCP state start time. • Last completed—IPCP state authentication completion time. • Negotiated options: <ul style="list-style-type: none"> • compression protocol—Negotiate the use of a specific compression protocol. By default, compression is not enabled. • local address—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. • primary DNS server—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link. • primary WINS server—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link. • remote address—IP address of the remote end of the link in dotted quad notation. • secondary DNS server—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link. • secondary WINS server—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: Active or Passive 	extensive

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPV6CP	<p>Internet Protocol version 6 Control Protocol (IPv6CP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—IPv6CP state start time. • Last completed—IPv6CP state authentication completion time. • Negotiated options: <ul style="list-style-type: none"> • local interface identifier—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. • remote interface identifier—IP address of the remote end of the link in dotted quad notation. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: Active or Passive 	extensive

Table 31: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSINLCP State	<p>OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):</p> <ul style="list-style-type: none"> • State: <ul style="list-style-type: none"> • Ack-rcvcd—Configure-Request has been sent and Configure-Ack has been received. • Ack-sent—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—Attempt has been made to configure the connection. • Last started—OSINLCP state start time. • Last completed—OSINLCP state completion time. 	extensive
TAGCP	<p>TAGCP information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvcd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—TAGCP state start time. • Last completed—TAGCP state authentication completion time. 	extensive none

Sample Output

show ppp interface

```
user@host> show ppp interface si-1/3/0.0
Session si-1/3/0.0, Type: PPP, Phase: Authenticate
Session flags: Monitored
LCP State: Opened
AUTHENTICATION: CHAP State: Chap-resp-sent, Chap-ans-sent
IPCP State: Closed, OSINLCP State: Closed
```

show ppp interface extensive

```
user@host> show ppp interface si-0/0/3.0 extensive
```

```
Session si-0/0/3.0, Type: PPP, Phase: Network
Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
                    Magic-Number validation: disable
RE Keepalive statistics:
LCP echo req Tx      : 657 (last sent 00:50:10 ago)
LCP echo req Rx      : 0 (last seen: never)
LCP echo rep Tx      : 0
LCP echo rep Rx      : 657
LCP echo req timeout : 0
LCP Rx echo req Magic Num Failures : 0
LCP Rx echo rep Magic Num Failures : 0
LCP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Authentication protocol: PAP, Magic number: 2341124815, MRU: 4470
Authentication: PAP
State: Success
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
IPCP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Local address: 203.0.113.21, Remote address: 203.0.113.22
Negotiation mode: Active
IPV6CP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Local interface identifier: 2a0:a522:64:d319, Remote interface identifier: 0:0:0:c
Negotiation mode: Passive
```

show ppp interface terse

```
user@host> show ppp interface si-1/3/0 terse
Session name  Session type  Session phase  Session flags
si-1/3/0.0    PPP           Authenticate    Monitored
```

show ppp statistics

Syntax show ppp statistics
 <detail>
 <memory>
 <recovery>

Release Information Command introduced in Junos OS Release 7.5.

Description Display PPP interface statistics information.

Options **detail**—(Optional) Display the detailed statistics.

memory—(Optional) Display PPP process memory statistics.

recovery—(Optional) Display recovery state of PPP after a GRES or restart. It is safe to force another GRES or restart only when the recovery state indicates the recovery is done.



NOTE: When you issue this command option during the recovery process, the command may time out or fail silently rather than display output. Recovery is not complete until the command displays **Recovery state: recovery done**.

Required Privilege Level view

List of Sample Output [show ppp statistics on page 720](#)
[show ppp statistics detail on page 720](#)
[show ppp statistics recovery \(Safe to Restart\) on page 721](#)
[show ppp statistics recovery \(Unsafe to Restart\) on page 721](#)

Output Fields [Table 32 on page 715](#) lists the output fields for the **show ppp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 32: show ppp statistics Output Fields

Field Name	Field Description	Level of Output
Total sessions	Number of PPP sessions on an interface.	none detail

Table 32: show ppp statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Sessions in disabled phase	Number of PPP sessions disabled. Number of sessions where the link is either administratively or physically down. Once the PPP process learns from the kernel that Layer 2 is ready to send and receive traffic, it will do a phase transition from disabled to established. When LCP and NCP transitions through states, links transition to the establish phase when terminate packets are exchanged or some other failure, such as authentication or expiration of a timer occurs.	none detail
Sessions in establish phase	Number of PPP sessions in establish phase. In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure and test the data link.	none detail
Sessions in authenticate phase	Number of PPP sessions in authenticate phase. Each end of the PPP link must first send LCP packets to configure the data link during the link establishment phase. After the link has been established, PPP provides for an optional authentication phase before proceeding to the Network-Layer Protocol (NLP) phase.	none detail
Sessions in network phase	Number of PPP sessions in the network phase. After a link has been established and optional facilities have been negotiated as needed by the LCP, PPP must send Network Control Protocol (NCP) packets to choose and configure one or more network-layer protocols, such as IP, IPX, or AppleTalk. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.	none detail
Bundles in pending phase	Number of unique bundles to which PPP links are referring.	none detail

Table 32: show ppp statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type	<p>Type of structure for which memory is allocated.</p> <ul style="list-style-type: none"> • Queued rtsock msgs—Queued route socket messages. When a PPP process is unable to send a route socket message to the kernel (typically because of congestion of the route socket interface), the message is queued for deferred processing. • PPP session—Active PPP session. Stores all the information for a PPP session, such as authentication, sequence number, LCP session, and NCP session information. • Interface address—Interface address associated with a PPP connection. Stores the information about the interface address that PPP obtains from the kernel. • Destination profile—Stores the destination profile information associated with an interface address. • ML link settings—Stores information about an MLPPP link, such as the bundle name and compressed real-time transport protocol (CRTP) settings. • IPCP blocked address—When addresses are blocked in an address pool (for example, when the interface address is within the range of an address pool, it will be implicitly blocked), this structure is used to store the address in the pool. • PPP session trace—A PPP session trace is allocated for record keeping for each session listed at the [set protocols ppp monitor-session] hierarchy level. • IFL redundancy state—Stores redundancy state information needed for high availability (HA) operation. • Protocol family—Stores the information about the protocol family that PPP obtains from the kernel. 	detail

Table 32: show ppp statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type (continued)	<ul style="list-style-type: none"> • ML bundle settings—Multilink bundle settings. Stores the context information for a MLPPP bundle. • PPP LCP session—PPP Link Control Protocol session, used for establishing, configuring, and testing the data-link connection. Stores the information for an LCP session, such as negotiated options, current state, and statistics. • PPP NCP session—PPP Network Control Protocol (NCP) phase in the PPP link connection process. Stores the information for an NCP session, such as negotiated options, current state, address family, and statistics. • Physical interface—Stores the information about the physical interface that PPP obtains from the kernel. • Access profile—Stores the information found at the [edit access profile] hierarchy level for each profile. • ML wait entry—Created when there are MLPPP links joining a bundle. before its addition to the PPP process. Links are saved here, and when the bundle is added, are properly assigned to the bundle. • Group profile—Stores information set in the PPP stanza of a group profile, such as the primary and secondary Domain Name System (DNS), primary and secondary NDNS, and address pool name. • Profile client—Stores the per-client information of the access profile (information obtained from the [set access profile name client client-name] hierarchy level. • PPP Auth session—PPP authentication session. Stores all the session-specific authentication protocol parameters. • Logical interface—Stores the information about the logical interface that PPP obtains from the kernel. • Non-tagged—Generic catch-all for allocations not of a particular structure type. 	detail

Table 32: show ppp statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type	<p>If you specify the memory keyword, the following memory statistics are displayed for Ethernet interfaces on M120 and M320 routers.</p> <ul style="list-style-type: none"> • authenticate—Stores information common to all PPP authentication protocols. • linkInterface—Stores information about PPP link interfaces. • pap—Stores information about PPP PAP authentication protocol. Includes authenticator and authenticate state machines. • lcp—PPP Link Control Protocol session. Used for establishing, configuring and testing the data-link connection. Stores information for LCP session, such as negotiated options, state, and statistics. • chap—Stores information about PPP CHAP authentication protocol. Includes authenticator and authenticate state machines. • eapBuffer—Stores runtime authentication information for EAP. • eap—Stores information about PPP EAP authentication protocol. Includes authenticator and authenticate state machines. • authNone—Stores information about no PPP authentication. Includes the authenticator state machine. • networkInterface—Stores information about NCP portions of PPP protocol. • ipNcp—PPP IPCP session information. Used for configuring, negotiating, and establishing IPCP protocol. Stores the current state, and configured and negotiated options. • ipv6Ncp—PPP IPv6CP session information. Used for configuring, negotiating, and establishing IPv6CP protocol. Stores the current state, and configured and negotiated options. • osiNcp—PPP OSICP session information. Used for configuring, negotiating, and establishing OSICP protocol. Stores the current state, and configured and negotiated options. • mplsNcp—PPP MPLSCP session information. Used for configuring, negotiating, and establishing MPLSCP protocol. Stores the current state. • trace—Stores information for PPP debugging. 	memory
Total	Total memory allocations.	detail
Size	Size of the structure.	detail
Active	Number of instances of the structure that are used.	detail
Free	Number of instances of the structure that are on the free list. Types with a number in the Free column are pooled structures, and are typically types that are often used.	detail
Limit	Maximum number of instances that can be on the free list. Types with a number in the Limit column are pooled structures, and are typically types that are often used.	detail
Total size	Total amount of memory being used by a type of structure (includes active and free instances).	detail
Requests	Number of allocation requests made by a type of structure.	detail

Table 32: show ppp statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Failures	Number of failed allocations.	detail
Recovery state	State of PPP recovery after a GRES or restart: <ul style="list-style-type: none"> recovery done—All sessions have recovered; it is safe to force another GRES or restart. recovery cleanup pending—Not all PPP sessions have recovered; it is not safe to force another GRES or restart. 	none
Subscriber sessions pending retention	Number of PPP subscriber sessions that are in the process of being recovered.	none
Subscriber sessions recovered OK	Number of PPP subscriber sessions that have recovered after a GRES or restart.	none
Subscriber sessions recovery failed	Number of PPP subscriber sessions that have failed to recover after a GRES or restart.	none

Sample Output

show ppp statistics

```

user@host> show ppp statistics
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0

```

show ppp statistics detail

```

user@host> show ppp statistics detail
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0
Type      Size  Active  Free  Limit  Total size  Requests  Failures
Queued rtsock msgs  28    0    0  65535      0        0
PPP session        60    0    0      0      0        0
Interface address  64    0    0  65535      0        0
Destination profile 65    0    0      0      0        0

```

ML link settings	68	0			0	0
IPCP blocked address	68	0			0	0
PPP session trace	76	0			0	0
IFL redundancy state	76	0			0	0
Protocol family	84	0	0	65535	0	0
ML bundle settings	108	0			0	0
PPP LCP session	120	0			0	0
PPP NCP session	124	0			0	0
Physical interface	124	170	0	65535	21080	170
Access profile	132	0			0	0
ML wait entry	144	0	0	20	0	0
Group profile	164	0			0	0
Profile client	272	0			0	0
PPP Auth session	356	0			0	0
Logical interface	524	0	0	65535	0	0
Non-tagged					8	2
Total					21088	172

0

Session statistics from PPP universal edge process

Total subscriber sessions: 32

Subscriber sessions in disabled phase : 32

Subscriber sessions in establish phase : 0

Subscriber sessions in authenticate phase: 0

Subscriber sessions in network phase : 0

Type	Size	Active	Free	Limit	Total	size	Requests	Failures
authenticate	224	1	99	16384	224	0	0	0
linkInterface	152	1	99	16384	152	0	0	0
pap	256	1	99	16384	256	0	0	0
lcp	272	1	99	16384	272	0	0	0
chap	284	0	0	16384	0	0	0	0
eapBuffer	1464	0	0	16384	0	0	0	0
eap	276	0	0	16384	0	0	0	0
authNone								
networkInterface	220	1	99	16384	220	0	0	0
ipNcp	256	1	99	16384	256	0	0	0
ipv6Ncp	204	0	0	16384	0	0	0	0
osiNcp	192	0	0	16384	0	0	0	0
mplsNcp	188	0	0	16384	0	0	0	0
trace	2052	0	16	16	0	0	0	0
Total					1380	0	0	0

show ppp statistics recovery (Safe to Restart)

user@host> show ppp statistics recovery

Recovery statistics from PPP universal edge process

Recovery state: recovery done

Subscriber sessions recovered OK : 32001

Subscriber sessions recovery failed : 0

show ppp statistics recovery (Unsafe to Restart)

user@host> show ppp statistics recovery

Recovery statistics from PPP universal edge process

Recovery state: recovery cleanup pending

Subscriber sessions pending retention : 32001

Subscriber sessions recovered OK : 0

Subscriber sessions recovery failed : 0

show ppp summary

Syntax	show ppp summary
Release Information	Command introduced in Junos OS Release 7.5.
Description	Display PPP session summary information.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show ppp summary on page 722
Output Fields	Table 33 on page 722 lists the output fields for the show ppp summary command. Output fields are listed in the approximate order in which they appear.

Table 33: show ppp summary Output Fields

Field Name	Field Description
Interface	Interface on which the PPP session is running. An interface type of pp0 indicates an Ethernet interface type on a M120 or M320 router.
Session type	Type of session: PPP or Cisco-HDLC .
Session phase	PPP process phases: Authenticate , Pending , Establish , Network , Disabled .
Session flags	Special conditions present in the session, such as Bundled , TCC , No-keepalives , Looped , Monitored , and NCP-only .

Sample Output

show ppp summary

```

user@host> show ppp summary
Interface    Session type  Session phase  Session flags
at-4/0/0.456 PPP           Network
lsq-0/3/0.0  PPP           Disabled
lsq-1/0/0.0  PPP           Disabled
r1sq0.0      PPP           Network        NCP-only
so-1/0/0.0   PPP           Authenticate
so-1/0/1.0   PPP           Disabled        Looped
so-2/0/0.0   Cisco-HDLC    Establish
so-4/0/0.0   PPP           Establish        Monitored
t1-1/3/0:1.0 PPP           Network        Bundled
t1-1/3/0:2.0 PPP           Network        Bundled
pp0.12       PPP           Network

```


show services inline ip-reassembly statistics

Syntax `show services inline ip-reassembly statistics`
 `<fpc fpc-slot>`
 `<pfe pfe-slot>`

Release Information Statement introduced in Junos OS Release 12.2X49.

Description Display the inline IP reassembly statistics for the Packet Forwarding Engines on one or more MPCs. Inline IP reassembly statistics are collected at the Packet Forwarding Engine level.



NOTE: For more information on MPCs that support inline IP reassembly, refer to *Protocols and Applications Supported on the MPC1E for MX Series Routers*.

Options **none**—Displays standard inline IP reassembly statistics for all MPCs.

fpc fpc—(Optional) Displays inline IP reassembly statistics for the specified MPC.



NOTE: Starting with Junos OS Release 14.2, the FPC option is not displayed for MX Series routers that do not contain switch fabrics, such as MX80 and MX104 routers.

pfe pfe—(Optional) Displays inline IP reassembly for the specified Packet Forwarding Engine slot. You must specify an FPC slot number before specifying a Packet Forwarding Engine slot.

Required Privilege Level view

Related Documentation • [ip-reassembly on page 514](#)

List of Sample Output [show services inline ip-reassembly statistics fpc on page 728](#)

Output Fields [Table 34 on page 725](#) lists the output fields for the `show services inline ip-reassembly statistics` command. Output fields are listed in the approximate order in which they appear.

Table 34: show services inline ip-reassembly statistics Output Fields

Field Name	Field Description
FPC	MPC slot number for which the statistics are displayed.
PFE	Packet Forwarding Engine on the MPC for which the statistics are displayed.

NOTE: The output fields displayed (per Packet Forwarding Engine) are arranged in a logical sequence from top to bottom to enable users to understand how the inline IP reassembly statistics are gathered.

The information about total number of fragments received is displayed first, and then the information about the reassembled packets and those pending reassembly are displayed. Then, the reasons why the fragments were dropped or not reassembled are displayed. Finally, the information about the fragments reassembled, fragments dropped, and fragments sent to the backup user plane PIC (services PIC) are displayed.

Total Fragments Received	<p>Total number of fragments received and the current rate of fragments received for inline IP reassembly. The following information is also displayed:</p> <ul style="list-style-type: none"> • First Fragments—Number of first fragments received and current rate of first fragments processed. • Intermediate Fragments—Number of intermediate fragments received and current rate of intermediate fragments processed. • Last Fragments—Number and rate of last fragments received. <p>NOTE: Current rate refers to the current number of fragments processed per second in the instant preceding the command's execution.</p>
Total Packets Reassembled	Total number of packets reassembled and current rate, in the instant preceding the command's execution, at which the packets are reassembled.
Approximate Packets Pending Reassembly	Approximate number of packets pending reassembly.

Table 34: show services inline ip-reassembly statistics Output Fields (*continued*)

Field Name	Field Description
Fragments Dropped Reasons	<p>Total number of fragments dropped reasons and the current rate of total fragment dropped reasons. The number of dropped reasons and rate corresponding to each of the following reasons are also displayed:</p> <ul style="list-style-type: none"> • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment <p>NOTE:</p> <ul style="list-style-type: none"> • These fields indicate <i>why</i> a fragment was dropped. When a fragment is dropped, the corresponding reason field is incremented by 1. For example, when a fragment is dropped because the memory runs out, the Buffers not available field increases by 1. • The maximum number of fragments allowed for reassembly is 16. If the interface encounters a 17th fragment, it drops the entire packet and increments the Fragment per packet exceeded field by 17. • Current rate refers to the current number of fragment dropped reasons per second in the instant preceding the command's execution.
Reassembly Errors Reasons	<p>Number of errors during reassembly and the current rate of reassembly errors. The number of errors and the rate for each of the following types of errors are also displayed:</p> <ul style="list-style-type: none"> • Fragment not found • Fragment not in sequence • ASIC errors <p>NOTE: Current rate refers to the current number of reassembly errors processed per second in the instant preceding the command's execution.</p>
Aged out packets	<p>Number of aged out packets and the current number of packets aged out per second in the instant preceding the command's execution.</p> <p>NOTE: In some cases, aged out packets can refer to aged out fragments. If previous fragments of the packet have already been discarded then linking of the dropped fragments to the aged out fragments cannot occur.</p>
Total Fragments Successfully Reassembled	<p>Number of fragments successfully reassembled and the current number of fragments reassembled per second in the instant preceding the command's execution.</p>

Table 34: show services inline ip-reassembly statistics Output Fields (*continued*)

Field Name	Field Description
Total Fragments Dropped	<p>Total number of fragments dropped and the current rate of total number of fragments dropped. The number of fragments dropped and rate corresponding to each of the following reasons are also displayed:</p> <ul style="list-style-type: none"> • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment • Fragment not found • Fragment not in sequence • ASIC errors • Aged out fragments
Total fragments punted to UPIC	Number of fragments sent to the backup user plane PIC (services PIC) and current rate of fragments sent per second in the instant preceding the command's execution

The following information applies to the **Total Fragments Dropped** field.

- These fields indicate *how many* of the packet fragments received were then dropped due to a particular reason.

For example, consider a packet that has 10 fragments, 9 of which have been received and stored in memory. When the tenth fragment arrives, if the memory runs out (Buffers not available), then this fragment is dropped. Because the tenth fragment has been dropped, the other 9 fragments must also be dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 10.

For the next packet arriving, which also has 10 fragments, the first four fragments are stored but the memory runs out for the fifth fragment. Then the first 5 fragments (fifth and the first four) are dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 5.

For the remaining fragments of the packet, if memory becomes available, the next 5 fragments (6 through 10) that arrive are stored in memory. The fragments are stored until the timeout period elapses, and are eventually dropped. In this case, the **Aged out packets** field is incremented by 1 and the **Aged out fragments** field (under the **Total Fragments Dropped** field) is incremented by 5.

The fragment counters (after both packets have been processed) are as follows:

- **Fragments Dropped Reasons**
 - Buffers not available 2
 - Aged out packets 1
- **Total Fragment Dropped**
 - Buffers not available 15
 - Aged out packets 5
- Current rate refers to the current total number fragments dropped per second in the instant preceding the command's execution.

Sample Output

show services inline ip-reassembly statistics fpc

```

user@host> show services inline ip-reassembly statistics fpc 0
FPC: 0 PFE: 0
=====

```

	Total	Current Rate
Total Fragments Received	728177644	83529
First Fragments	260759430	29924
Intermediate Fragments	206658784	23681
Last Fragments	260759430	29924
Total Packets Successfully Reassembled	260746982	29924
Approximate Packets Pending Reassembly	4	
Fragments Dropped Reasons	34558	3
Buffers not available	0	0
Fragments per packet exceeded	0	0
Packet length exceeded	0	0
Record insert error	0	0
Record in use error	34558	3
Duplicate first fragments	0	0
Duplicate last fragments	0	0
Missing first fragment	0	0
Reassembly Errors Reasons	0	0
Fragment not found	0	0
Fragment not in sequence	0	0
ASIC errors	0	0
Aged out packets	63	0
Total Fragments Successfully Reassembled	728142977	83528
Total Fragments Dropped	34673	3
Buffers not available	0	0
Fragments per packet exceeded	0	0
Packet length exceeded	0	0
Record insert error	0	0
Record in use error	34558	3
Duplicate first fragments	0	0
Duplicate last fragments	0	0
Missing first fragment	0	0

Fragment not found	0	0
Fragment not in sequence	0	0
ASIC errors	0	0
Aged out fragments	115	0
Total fragments punted to UPIC	0	0

show services l2tp client

Syntax	<code>show services l2tp client</code> <code><client-name></code>
Release Information	Command introduced in Junos OS Release 16.1.
Description	Display information about all L2TP clients or a specific L2TP client.
Options	client-name —(Optional) Name of a client.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp session-limit-group on page 747 • show services l2tp tunnel-group on page 761 • L2TP Session Limits Overview on page 199
List of Sample Output	show services l2tp client on page 730 show services l2tp client (Client Name) on page 731
Output Fields	Table 35 on page 730 lists the output fields for the show services l2tp client command. Output fields are listed in the approximate order in which they appear.

Table 35: show services l2tp client Output Fields

Field Name	Field Description
Client	Name of the client.
Client Name	
Tunnels	Number of tunnels in the tunnel group.
Sessions	Number of L2TP sessions established for tunnels in the tunnel group.
Tunnel-group	Name of a tunnel group to which the client belongs.
Session-limit-group	Name of a session-limit group to which the client belongs.

Sample Output

show services l2tp client

```
user@host> show services l2tp client
```

Client	Tunnels	Sessions	Tunnel-group	Session-limit-group
entA-serviceA	2	20	l2tp-tunnel-group1	enterpriseA
entA-serviceB	3	120	l2tp-tunnel-group2	enterpriseB

show services l2tp client (Client Name)

```
user@host> show services l2tp client entA-serviceA
```

Client Name	Tunnels	Sessions	Tunnel-group	Session-limit-group
entA-serviceA	2	20	l2tp-tunnel-group1	enterpriseA

show services l2tp destination

Syntax	<code>show services l2tp destination</code> <code><brief detail extensive></code> <code><local-gateway <i>gateway-address</i>></code> <code><peer-gateway <i>gateway-address</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display information about L2TP tunnel destinations.
Options	<p>brief detail extensive—(Optional) Display the specified level of information.</p> <p>local-gateway <i>gateway-address</i>—(Optional) Display L2TP session information for only the specified local gateway address.</p> <p>peer-gateway <i>gateway-address</i>—(Optional) Display L2TP session information for only the specified peer gateway address.</p> <p>statistics—(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with any of the level options, brief, detail, or extensive.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear services l2tp destination on page 663• show services l2tp destination lockout on page 736• show services l2tp session on page 738• show services l2tp summary on page 749• show services l2tp tunnel on page 755
List of Sample Output	show services l2tp destination on page 734 show services l2tp destination detail on page 734 show services l2tp destination extensive (LAC) on page 735 show services l2tp destination extensive (LNS) on page 735 show services l2tp destination statistics (LAC only on MX Series Routers) on page 735
Output Fields	Table 36 on page 733 lists the output fields for the show services l2tp destination command. Output fields are listed in the approximate order in which they appear.

Table 36: show services l2tp destination Output Fields

Field Name	Field Description	Level of Output
Local Name	Name of this destination.	All levels
Remote IP	IP address of the remote peer (LNS).	All levels
Tunnels	Number of tunnel connections for the destination in the following categories: <ul style="list-style-type: none"> total active failed 	All levels for total extensive for active and failed
Sessions	Number of session connections for the destination in the following categories: <ul style="list-style-type: none"> total active failed 	All levels for total extensive for active and failed
State	Administrative state of the L2TP destination: <ul style="list-style-type: none"> Enabled—No restrictions exist on creation or operation of sessions and tunnels for this destination. Disabled—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the Disabled state. Drain—Creation of new sessions and tunnels is disabled for this destination. 	All levels
Local IP	IP address of the local gateway (LAC).	detail extensive
Transport	Medium used for tunneling. Only ipUdp is supported.	detail extensive
Logical System	Logical system in which the tunnel is configured.	detail extensive
Router Instance	Routing instance in which the tunnel is configured.	detail extensive
Lockout State	Reachability state of the destination: <ul style="list-style-type: none"> not locked—Destination is considered reachable. waiting for lockout timeout—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber. 	detail extensive
Access Line Information	State of the LAC per-destination configuration for forwarding subscriber line information to the LNS, Enabled or Disabled . Starting in Junos OS Release 17.4R1, this information is displayed on the LNS for information it receives from the LAC, Enabled or Disabled .	detail extensive

Table 36: show services l2tp destination Output Fields (*continued*)

Field Name	Field Description	Level of Output
Speed Updates	State of the LAC per-destination configuration for including connection speed updates when it forwards subscriber line information to the LNS, Enabled or Disabled . Starting in Junos OS Release 17.4R1, this information is displayed on the LNS for updates it receives from the LAC, Enabled or Disabled .	detail extensive
Connections	Number of total, active, and failed tunnel and session connections for the destination.	extensive
Control Tx	Amount of control information transmitted, in packets and bytes.	statistics
Control Rx	Amount of control information received, in packets and bytes.	statistics
Data Tx	Amount of data transmitted, in packets and bytes.	statistics
Data Rx	Amount of data received, in packets and bytes.	statistics
Error Tx	Number of errors transmitted, in packets.	statistics
Error Rx	Number of errors received, in packets.	statistics

Sample Output

show services l2tp destination

```

user@host> show services l2tp destination
  Local Name  Remote IP    Tunnels    Sessions  State
  1           203.0.113.101  1          1         Enabled

```

show services l2tp destination detail

```

user@host> show services l2tp destination detail
Local name: 1
  Remote IP: 203.0.113.101
  Tunnels: 1, Sessions: 1
  State: Enabled
  Local IP: 203.0.113.102
  Transport: ipUdp, Logical System: default, Router Instance: default
  Lockout State: not locked
  Access Line Information: Enabled, Speed Updates: Enabled
Local name: 1
  Remote IP: 203.0.113.108
  Tunnels: 1, Sessions: 1
  State: Enabled
  Local IP: 203.0.113.2
  Transport: ipUdp, Logical System: default, Router Instance: default
  Lockout State: waiting for lockout timeout
  Access Line Information: Enabled, Speed Updates: Enabled

```

show services l2tp destination extensive (LAC)

```

user@host> show services l2tp destination extensive
Local name: 1
Remote IP: 203.0.113.101
State: Enabled
Local IP: 203.0.113.102
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Access Line Information: Enabled, Speed Updates: Enabled
Connections      Totals      Active      Failed
Tunnels          1           1           0
Sessions         1           1           0

```

show services l2tp destination extensive (LNS)

```

user@host> show services l2tp destination extensive
Local name: 3
Remote IP: 203.0.113.103
State: Enabled
Local IP: 203.0.113.102
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Access Line Information: Enabled, Speed Updates: Disabled
Connections      Totals      Active      Failed
Tunnels          1           1           0
Sessions         1           1           0

```

show services l2tp destination statistics (LAC only on MX Series Routers)

```

user@host> show services l2tp destination statistics
Local name: 2, Tunnels: 1, Sessions: 210

```

	Packets	Bytes
Control Tx	680	63.3k
Control Rx	283	10.6k
Data Tx	1129	14.3k
Data Rx	877	10.9k
Errors Tx	0	
Errors Rx	0	

show services l2tp destination lockout

Syntax	show services l2tp destination lockout
Release Information	Command introduced in Junos OS Release 13.2.
Description	Display a list of destinations that are currently locked out and the time remaining for each to remain in the lockout state.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear services l2tp destination on page 663 • request services l2tp destination unlock on page 595 • show services l2tp destination on page 732 • show services l2tp session on page 738 • show services l2tp summary on page 749 • show services l2tp tunnel on page 755
List of Sample Output	show services l2tp destination lockout on page 736
Output Fields	Table 37 on page 736 lists the output fields for the show services l2tp destination lockout command. Output fields are listed in the approximate order in which they appear.

Table 37: show services l2tp destination lockout Output Fields

Field Name	Field Description
Destination	Name of the destination.
Time Remaining	Time remaining for the destination to be locked out.
L2TP lockout destinations found	Total count of lockout destinations.

Sample Output

show services l2tp destination lockout

```

user@host> show services l2tp destination lockout
  Destination  Time Remaining
  ---
  4             45
  5             43

```

```
6           8
3 L2TP lockout destinations found
```

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](#)
- [L2TP Minimum Configuration](#)
- [clear services l2tp session on page 667](#)

List of Sample Output

- [show services l2tp session \(LNS on M Series Routers\) on page 743](#)
- [show services l2tp session \(LNS on MX Series Routers\) on page 743](#)
- [show services l2tp session \(LAC\) on page 743](#)
- [show services l2tp session detail \(LAC\) on page 743](#)
- [show services l2tp session extensive \(LAC\) on page 744](#)
- [show services l2tp session extensive \(LAC on MX Series Routers\) on page 744](#)
- [show services l2tp session extensive \(LNS on M Series Routers\) on page 744](#)
- [show services l2tp session extensive \(LNS on MX Series Routers\) on page 745](#)
- [show services l2tp session statistics \(MX Series Routers\) on page 745](#)

Output Fields [Table 38 on page 739](#) lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

Table 38: 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

Table 38: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	All levels
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
State	<p>State of the L2TP session:</p> <ul style="list-style-type: none"> • Established—Session is operating. This is the only state supported for the LAC. • closed—Session is being closed. • destroyed—Session is being destroyed. • clean-up—Session is being cleaned up. • lns-ic-accept-new—New session is being accepted. • lns-ic-idle—Session has been created and is idle. • lns-ic-reject-new—New session is being rejected. • lns-ic-wait-connect—Session is waiting for the peer's incoming call connected (ICCN) message. 	All levels
Bundle ID	(LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank Bundle field are not participating in the Multilink Protocol. Sessions in a multilink bundle might belong to different L2TP tunnels. For L2TP output organized by bundle ID, issue the show services l2tp multilink extensive command.	All levels
Mode	<p>(LNS) Mode of the interface representing the session: shared or exclusive.</p> <p>(LAC) Mode of the interface representing the session: shared or dedicated. Only dedicated is currently supported for the LAC.</p>	extensive
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	extensive
Remote IP	IP address of remote endpoint of the PPP session.	extensive
Username	(LNS only) Name of the user logged in to the session.	All levels
Assigned IP address	(LNS only) IP address assigned to remote client.	extensive
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	extensive
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	extensive
Local MRU	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	extensive
Remote MRU	(LNS only) MRU setting of the remote device, in bytes.	extensive

Table 38: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Tx speed	<p>Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived.</p> <p>Starting in Junos OS Release 14.1, either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>For Junos OS Release 17.2 and Release 17.3, only the current (update) line speed can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 17.4R1, once again either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	extensive
Rx speed	<p>Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived.</p> <p>Starting in Junos OS Release 14.1, either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>For Junos OS Release 17.2 and Release 17.3, only the current (update) line speed can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 17.4R1, once again either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	extensive
Bearer type	<p>Type of bearer enabled:</p> <ul style="list-style-type: none"> • 0—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • 1—Digital access requested. • 2—Analog access requested. • 4—Asynchronous Transfer Mode (ATM) bearer support. 	extensive
Framing type	<p>Type of framing enabled:</p> <ul style="list-style-type: none"> • 1—Synchronous framing • 2—Asynchronous framing 	extensive

Table 38: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
LCP renegotiation	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: On or Off .	extensive
Authentication	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	extensive
Interface ID	(LNS only) Identifier used to look up the logical interface for this session.	extensive
Interface unit	Logical interface for this session.	All levels
Call serial number	Unique serial number assigned to the call.	extensive
Policer bandwidth	Maximum policer bandwidth configured for this session.	extensive
Policer burst size	Maximum policer burst size configured for this session.	extensive
Firewall filter	Configured firewall filter name.	extensive
Session encapsulation overhead	Overhead allowance configured for this session, in bytes.	extensive
Session cell overhead	Cell overhead activation (On or Off).	extensive
Create time	Date and time when the call was created.	extensive
Up time	Length of time elapsed since the call became active, in hours, minutes, and seconds.	extensive
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	extensive

Table 38: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics since	Date and time when collection of the following statistics began: <ul style="list-style-type: none"> Control Tx—Amount of control information transmitted, in packets and bytes. Control Rx—Amount of control information received, in packets and bytes. Data Tx—Amount of data transmitted, in packets and bytes. Data Rx—Amount of data received, in packets and bytes. Errors Tx—Number of errors transmitted, in packets. Errors Rx—Number of errors received, in packets. LCP echo req Tx—Number of LCP echo requests transmitted, in packets. LCP echo req Rx—Number of LCP echo requests received, in packets. LCP echo rep Tx—Number of LCP echo responses transmitted, in packets. LCP echo rep Rx—Number of LCP echo responses received, in packets. LCP echo Req timeout—Number of LCP echo requests that timed out. LCP echo Req error—Number of errors received for LCP echo packets. LCP echo Rep error—Number of errors transmitted for LCP echo packets. 	extensive

Sample Output

show services l2tp session (LNS on M Series Routers)

```

user@host> show services l2tp session
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 8802
  Local Remote Interface State      Bundle Username
  ID   ID   unit
37966    5       2 Established

```

show services l2tp session (LNS on MX Series Routers)

```

user@host> show services l2tp session
Tunnel local ID: 40553
  Local Remote State      Interface      Interface
  ID   ID              unit          Name
17967  1      Established  1073749824  si-5/2/0

```

show services l2tp session (LAC)

```

user@host> show services l2tp session
Tunnel local ID: 31889
  Local Remote State      Interface      Interface
  ID   ID              unit          Name
31694    1      Established    311          pp0

```

show services l2tp session detail (LAC)

```

user@host> show services l2tp session detail
Tunnel local ID: 31889
Session local ID: 31694, Session remote ID: 1, Interface unit: 311
State: Established, Interface: pp0, Mode: Dedicated
Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
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: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    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: 203.0.113.102:1701, Remote IP: 203.0.113.101:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 256000, source service-profile
    Rx speed: 128000, source ancp
    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: 203.0.113.121:1701, Remote IP: 203.0.113.202:1701
    Username: user@example.com, Assigned IP address: 203.0.113.51/32
    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: 203.0.113.121:1701, Remote IP: 203.0.113.222:1701
Username: usr1@company.example.com, Assigned IP address: 203.0.113.3/24
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: 192.0.2.2:1701, Remote IP: 192.0.2.3:1701
Local name: lns-mx960, Remote name: testlac
Tx speed: initial 64000, Update 256000
Rx speed: initial 64000, Update 256000
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 session-limit-group

Syntax	show services l2tp session-limit-group <limit-group-name>
Release Information	Command introduced in Junos OS Release 16.1.
Description	Display information about all session-limit groups or a specific session limit group.
Options	limit-group-name —(Optional) Name of a session-limit group.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp client on page 730 • show services l2tp tunnel-group on page 761 • L2TP Session Limits Overview on page 199
List of Sample Output	show services l2tp session-limit-group on page 747 show services l2tp session-limit-group (Limit Group Name) on page 748
Output Fields	Table 39 on page 747 lists the output fields for the show services l2tp session-limit-group command. Output fields are listed in the approximate order in which they appear.

Table 39: show services l2tp session-limit-group Output Fields

Field Name	Field Description
Session-limit-group	Name of a session-limit group.
Tunnels	Number of tunnels associated with the session-limit group in the tunnel group.
Sessions	Number of L2TP sessions established for session-limit group.
Maximum limit	Maximum number of sessions allowed for the session-limit group.

Sample Output

show services l2tp session-limit-group

```
user@host> show services l2tp session-limit-group
```

Session-limit-group	Tunnels	Sessions	
enterpriseA	2	10	1000
enterpriseB	10	120	2000

show services l2tp session-limit-group (Limit Group Name)

user@host> show services l2tp session-limit-group enterpriseA

Session-limit-group	Tunnels	Sessions	Maximum limit
enterpriseA	2	10	1000
enterpriseC	10	120	2000

show services l2tp summary

Syntax	show services l2tp summary <interface sp-fpc/pic/port> <statistics>
Release Information	Command introduced before Junos OS Release 7.4. Support for LAC on MX Series routers introduced in Junos OS Release 10.4. Support for LNS on MX Series routers introduced in Junos OS Release 11.4. Support for statistics option introduced in Junos OS Release 13.1.
Description	(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Display Layer 2 Tunneling Protocol (L2TP) summary information.
Options	<p>none—Display complete L2TP summary information. For LNS on M Series routers, display L2TP summary information for all adaptive services interfaces. For LNS on MX Series routers, display L2TP summary information for all inline services interfaces.</p> <p>interface sp-fpc/pic/port—(Optional) Display L2TP summary information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.</p> <p>statistics—(Optional) Display a summary of control packets and bytes transmitted and received.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>L2TP Services Configuration Overview</i> • <i>L2TP Minimum Configuration</i>
List of Sample Output	show services l2tp summary (LAC on M Series routers) on page 752 show services l2tp summary (LAC on MX Series routers) on page 753 show services l2tp summary (LNS on MX Series routers) on page 753 show services l2tp summary (LNS on M Series routers) on page 753 show services l2tp summary statistics (MX Series routers) on page 753
Output Fields	Table 40 on page 749 lists the output fields for the show services l2tp summary command. Output fields are listed in the approximate order in which they appear.

Table 40: show services l2tp summary Output Fields

Field Name	Field Description
Administrative state	Administrative state of the tunnel is drain. In this state you cannot configure new sessions, destinations, or tunnels at the LAC or LNS.

Table 40: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Failover within a preference level	State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level. Not displayed for LNS on M Series routers.
Weighted load balancing	State of this tunnel selection method on the LAC. When enabled, the maximum session limit of a tunnel determines its weight within a preference level. Tunnel selection proceeds from greatest to least weight. When disabled, selection defaults to a round robin method. Not displayed for LNS on M Series routers.
Destination equal load balancing	State of this tunnel selection method on the LAC. When enabled, the LAC selects tunnels based on the session count for destinations and the tunnel session count. Not displayed for LNS on M Series routers.
Tunnel authentication challenge	State of tunnel authentication, indicating whether the LAC and LNS exchange an authentication challenge and response during the establishment of the tunnel. The state is Enabled when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is Disabled when the secret is not present. Not displayed for LNS on M Series routers.
Calling number avp	When the state is Enabled , the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is Disabled , the attribute is not sent to the LNS. Not displayed for LNS on M Series routers.
Failover Protocol	When the state is enabled, the LAC operates in the default <i>failover-protocol-fall-back-to-silent-failover</i> manner. When the state is disabled, the disable-failover-protocol statement has been issued and the LAC operates only in silent failover mode. Not displayed for LNS on M Series routers.
Tx connect speed method	<p>The connection speed method configured to send the speed values in the L2TP Tx Connect Speed (AVP 24) and L2TP Rx Connect Speed (AVP 38). Possible values are:</p> <ul style="list-style-type: none"> • actual This is the default value in Junos OS Releases 15.1, 16.1, 16.2, and 17.1. It is deprecated in Junos Releases 17.2 and higher. • ancp • none • pppoe-ia-tag • service-profile • static This is the default value in Junos Releases 13.3, 14.1, 14.2, 17.2 and higher. It is deprecated in Junos OS Releases 15.1, 16.1, 16.2, and 17.1.
Rx speed avp when equal	Indicates if the Rx connect speed when equal configuration is enabled or disabled .

Table 40: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Tunnel assignment id	<p>Format of the tunnel name.</p> <p>Format of the tunnel name, based on RADIUS attributes returned from the AAA server:</p> <ul style="list-style-type: none"> • authentication-id—Name consists of only Tunnel Assignment-Id [82]. This is the default value. • client-server-id—Name is a combination of Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. This format is available only on MX Series routers.
Tunnel Tx Address Change	<p>Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both:</p> <ul style="list-style-type: none"> • accept—Accepts change requests for the IP address or UDP port. This is the default action. • ignore—Ignores all change requests. • ignore-ip-address—Ignores change requests for the IP address but accepts them for the UDP port. • ignore-udp-port—Ignores change requests for the UDP port but accepts them for the IP address.
Min Retransmission Timeout for control packets	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
Min Retransmission Timeout for control packets	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
Max Retransmissions for Established Tunnel	Maximum number of times control messages are retransmitted for established tunnels.
Max Retransmissions for Not Established Tunnel	Maximum number of times control messages are retransmitted for tunnels that are not established.
Tunnel Idle Timeout	Period that a tunnel can be inactive—that is, carrying no traffic—before it times out and is torn down.
Destruct Timeout	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
Reassembly Service Set	Indicates active IP reassembly configured for the interface.
Destination Lockout Timeout	Timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created.

Table 40: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Access Line Information	<p>State of LAC global configuration for forwarding subscriber line information to the LNS, Enabled or Disabled.</p> <p>Indicates active IP reassembly configured for the interface.</p> <p>Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for information it receives from the LAC.</p>
IPv6 Services for LAC Sessions	<p>State of LAC IPv6 service configuration for creating the IPv6 (inet6) address family for LAC subscribers, allowing the application of IPv6 firewall filters, Enabled or Disabled.</p>
Speed Updates	<p>State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, Enabled or Disabled.</p> <p>Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for updates it receives from the LAC.</p>
Destinations	Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers.
Tunnels	Number of L2TP tunnels established on the router.
Sessions	Number of L2TP sessions established on the router.
Switched sessions	Number of L2TP tunnel-switched sessions established on the router.
Control	Count of L2TP control packets and bytes sent and received.
Data	Count of L2TP data packets and bytes sent and received.
Errors	Count of L2TP error packets and bytes sent and received.

Sample Output

show services l2tp summary (LAC on M Series routers)

```

user@host> show services l2tp summary
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Enabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Disabled
Tunnel assignment id format is authentication-id
Destinations: 1 Tunnels: 1, Sessions: 1
  Tx packets    Rx packets  Memory (bytes)
Control    260           144           11513856

```

Data	7.5k	16.9k	8.3k
Errors	0	0	

show services l2tp summary (LAC on MX Series routers)

```
user@host> show services l2tp summary
Administrative state is Drain
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Enabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is enabled
  Tunnel Tx Address Change is Accept
  Min Retransmissions Timeout for control packets is 2 seconds
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
  Reassembly Service Set is ssnr3
  Access Line Information is Enabled, Speed Updates is Enabled
  IPv6 Services For LAC Sessions is Enabled
  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
  Access Line Information is Enabled, Speed Updates is Enabled
```

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 seconds
Destinations: 1, Tunnels: 1, Sessions: 31815, Switched sessions: 0

	Tx packets	Rx packets	Memory (bytes)	
Control		90.4k	32.0k	245678080
Data		127.3k	100.8kk	0
Errors		0	0	

show services l2tp tunnel

Syntax `show services l2tp tunnel`
 `<brief | detail | extensive>`
 `<interface sp-fpc/pic/port>`
 `<local-gateway gateway-address>`
 `<local-gateway-name gateway-name>`
 `<local-tunnel-id tunnel-id>`
 `<peer-gateway gateway-address>`
 `<peer-gateway-name gateway-name>`
 `<statistics>`
 `<tunnel-group group-name>`

Release Information Command introduced before Junos OS Release 7.4.

Description (M10i and M7i routers only) Display information about active Layer 2 Tunneling Protocol (L2TP) tunnels for LNS.

(MX Series routers only) Display information about L2TP tunnels for LAC and LNS; the tunnels may or may not have active sessions.

Options **none**—Display standard information about all active L2TP tunnels.

brief | detail | extensive—(Default) Display the specified level of output.

interface sp-fpc/pic/port—(Optional) Display L2TP tunnel information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.

local-gateway gateway-address—(Optional) Display L2TP tunnel information for only the specified local gateway address.

local-gateway-name gateway-name—(Optional) Display L2TP tunnel information for only the specified local gateway name.

local-tunnel-id tunnel-id—(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.

peer-gateway gateway-address—(Optional) Display L2TP tunnel information for only the specified peer gateway address.

peer-gateway-name gateway-name—(Optional) Display L2TP tunnel information for only the specified peer gateway name.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. The statistics for a tunnel are retained until the tunnel is disconnected, rather than until the last session in the tunnel is cleared. Retaining the statistics enables them to increment in the event a new session subsequently uses the tunnel. You cannot include this option with any of the level options, **brief**, **detail**, or **extensive**.

tunnel-group *group-name*—(Optional) Display L2TP tunnel information for only the specified tunnel group.

Required Privilege Level view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*

List of Sample Output

[show services l2tp tunnel \(LAC\) on page 758](#)
[show services l2tp tunnel detail \(LAC\) on page 758](#)
[show services l2tp tunnel detail \(LAC on MX Series Routers\) on page 758](#)
[show services l2tp tunnel detail \(LNS on MX Series Routers\) on page 759](#)
[show services l2tp tunnel extensive \(LAC\) on page 759](#)
[show services l2tp tunnel extensive \(LNS on M Series Routers\) on page 759](#)
[show services l2tp tunnel extensive \(LNS on MX Series Routers\) on page 760](#)
[show services l2tp tunnel statistics \(MX Series Routers\) on page 760](#)

Output Fields [Table 41 on page 756](#) lists the output fields for the **show services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 41: show services l2tp tunnel Output Fields

Field Name	Field Description
Interface	(LNS only) Name of an adaptive services interface.
Tunnel group	(LNS only) Name of a tunnel group.
Local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.
Remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.
Remote IP	IP address of the peer endpoint of the tunnel.
Sessions	Number of L2TP sessions established through the tunnel.

Table 41: show services l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
State	<p>State of the L2TP tunnel:</p> <ul style="list-style-type: none"> • cc_responder_accept_new—The tunnel has received and accepted the start control connection request (SCCRQ). • cc_responder_reject_new—The tunnel has received and rejected the SCCRQ. • cc_responder_idle—The tunnel has just been created. • cc_responder_wait_ctl_conn—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • clean-up—The tunnel is being cleaned up. • closed—The tunnel is being closed. • destroyed—The tunnel is being destroyed. • Drain—Creation of new sessions and destinations is disabled for this tunnel. • Established—The tunnel is operating. This is the only state supported for the LAC. • Terminate—The tunnel is terminating. • Unknown—The tunnel is not connected to the router.
Tunnel Name	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].
Local IP	IP address of the local endpoint of the tunnel.
Local name	Name used for local tunnel endpoint during tunnel negotiation.
Remote name	Name used for remote tunnel endpoint during tunnel negotiation.
Effective Peer Resync Mechanism	<p>(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel:</p> <ul style="list-style-type: none"> • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.
Nas Port Method	<p>NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS:</p> <ul style="list-style-type: none"> • cisco-avp—sends the AVP. • none—does not send the AVP.
Tunnel Logical System	Logical system in which the L2TP tunnel is brought up.
Tunnel Routing Instance	Routing instance in which the L2TP tunnel is brought up.
Max sessions	Maximum number of sessions that can be established on this tunnel.
Window size	Number of control messages that can be sent without receipt of an acknowledgment.
Hello interval	Interval between the transmission of hello messages, in seconds.

Table 41: show services l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
Create time	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to Unknown and the Create time value resets.
Up time	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.
Idle time	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.
Statistics since	<p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> • Control Tx—Amount of control information transmitted, in packets and bytes. • Control Rx—Amount of control information received, in packets and bytes. • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. • Errors Tx—Number of errors transmitted, in packets. • Errors Rx—Number of errors received, in packets.

Sample Output

show services l2tp tunnel (LAC)

```
user@host> show services l2tp tunnel
Local ID  Remote ID  Remote IP                Sessions  State
17185      1    203.0.113.101:1701      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: 192.0.2.1:1701
Sessions: 1, State: Established
Tunnel Name: 1/tunnel-to-LNS-1
Local IP: 192.0.2.2:1701
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: 203.0.113.101:1701
Sessions: 1, State: Established
Tunnel Name: 2/tunnel-to-LNS-2
Local IP: 192.0.2.2:1701
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: 198.51.100.15:1701
Sessions: 1, State: Established
Tunnel Name: 2/2
Local IP: 198.51.100.5:1701
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: 203.0.113.101:1701
Sessions: 1, State: Established
Tunnel Name: 2/tunnel-to-LNS-2
Local IP: 192.0.2.22:1701
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: 203.0.113.202:1701
Sessions: 1, State: Established
Local IP: 203.0.113.121:1701
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: 203.0.113.222:1701
Sessions: 1, State: Established
Local IP: 203.0.113.111:1701
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: 203.0.113.2:1701
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 tunnel-group

Syntax	<code>show services l2tp tunnel-group</code> <code><group-name></code>
Release Information	Command introduced in Junos OS Release 16.1.
Description	Display information about all L2TP tunnel groups or a specific L2TP tunnel group.
Options	<i>group-name</i> —(Optional) Name of a tunnel group.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp client on page 730 • show services l2tp session-limit-group on page 747 • L2TP Session Limits Overview on page 199
List of Sample Output	show services l2tp tunnel-group on page 761 show services l2tp tunnel-group (Group Name) on page 761
Output Fields	Table 42 on page 761 lists the output fields for the show services l2tp tunnel-group command. Output fields are listed in the approximate order in which they appear.

Table 42: show services l2tp tunnel-group Output Fields

Field Name	Field Description
Tunnel-group	Name of a tunnel group.
Tunnels	Number of tunnels in the tunnel group.
Sessions	Number of L2TP sessions established for tunnels in the tunnel group.

Sample Output

show services l2tp tunnel-group

```

user@host> show services l2tp tunnel-group
Tunnel-group      Tunnels      Sessions
l2tp-tunnel-group1  2             20
l2tp-tunnel-group2  3             120

```

show services l2tp tunnel-group (Group Name)

```

user@host> show services l2tp tunnel-group l2tp-tunnel-group1

```

Tunnel-group	Tunnels	Sessions
l2tp-tunnel-group1	2	20

show services l2tp tunnel-switch destination

Syntax	show services l2tp tunnel-switch destination < detail extensive > <statistics>
Release Information	Command introduced in Junos OS Release 13.2.
Description	Display information about L2TP switched tunnel destinations.
Options	<p>none—Display standard information for all L2TP switched tunnel destinations.</p> <p>detail extensive—(Optional) Display the specified level of information.</p> <p>statistics—(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with either of the level options, detail or extensive.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp tunnel-switch session on page 767 • show services l2tp tunnel-switch summary on page 772 • show services l2tp tunnel-switch tunnel on page 774
List of Sample Output	show services l2tp tunnel-switch destination on page 765 show services l2tp tunnel-switch destination detail on page 765 show services l2tp tunnel-switch destination extensive on page 765 show services l2tp tunnel-switch destination statistics on page 766
Output Fields	Table 43 on page 763 lists the output fields for the show services l2tp tunnel-switch destination command. Output fields are listed in the approximate order in which they appear.

Table 43: show services l2tp tunnel-switch destination Output Fields

Field Name	Field Description	Level of Output
Local Name	Name of this destination.	All levels
Remote IP	IP address of the remote peer (LNS).	All levels

Table 43: show services l2tp tunnel-switch destination Output Fields (*continued*)

Field Name	Field Description	Level of Output
Tunnels	Number of tunnel connections for the destination in the following categories: <ul style="list-style-type: none"> total active failed 	All levels for total extensive for active and failed
Sessions	Number of session connections for the destination in the following categories: <ul style="list-style-type: none"> total active failed 	All levels for total extensive for active and failed
Switched-sessions	Number of L2TP sessions established by tunnel switching.	All levels
State	Administrative state of the L2TP destination: <ul style="list-style-type: none"> Enabled—No restrictions exist on creation or operation of sessions and tunnels for this destination. Disabled—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the Disabled state. 	All levels
Local IP	IP address of the local gateway (LAC).	detail extensive
Transport	Medium used for tunneling. Only ipUdp is supported.	detail extensive
Logical System	Logical system in which the tunnel is configured.	detail extensive
Router Instance	Routing instance in which the tunnel is configured.	detail extensive
Lockout State	Reachability state of the destination: <ul style="list-style-type: none"> not locked—Destination is considered reachable. waiting for lockout timeout—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber. 	detail extensive
Connections	Number of total, active, and failed tunnel and session connections for the destination.	extensive
Control Tx	Amount of control information transmitted, in packets and bytes.	extensive statistics
Control Rx	Amount of control information received, in packets and bytes.	extensive statistics
Data Tx	Amount of data transmitted, in packets and bytes.	extensive statistics
Data Rx	Amount of data received, in packets and bytes.	extensive statistics

Table 43: show services l2tp tunnel-switch destination Output Fields (*continued*)

Field Name	Field Description	Level of Output
Error Tx	Number of errors transmitted, in packets.	extensive statistics
Error Rx	Number of errors received, in packets.	extensive statistics

Sample Output

show services l2tp tunnel-switch destination

```
user@host> show services l2tp tunnel-switch destination
Local Name Remote IP Tunnels Sessions Switched-sessions State
1          192.168.20.3 1      1      1           Enabled
2          203.0.113.10 1      1      1           Enabled
```

show services l2tp tunnel-switch destination detail

```
user@host> show services l2tp tunnel-switch destination detail
Local name: 1
Remote IP: 192.168.20.3
Tunnels: 1, Sessions: 1, Switched sessions: 1
State: Enabled
Local IP: 203.0.113.51
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Local name: 2
Remote IP: 198.51.100.10
Tunnels: 1, Sessions: 1, Switched sessions: 1
State: Enabled
Local IP: 203.0.113.31
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
```

show services l2tp tunnel-switch destination extensive

```
user@host> show services l2tp tunnel-switch destination extensive
Waiting for statistics...
Local name: 1
Remote IP: 192.168.20.3
Tunnels: 1, Sessions: 1, Switched sessions: 1
State: Enabled
Local IP: 203.0.113.51
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
Connections Totals Active Failed
Tunnels      1      1      0
Sessions     1      1      0
              Packets Bytes
Control Tx    6     239
Control Rx    6     267
Data Tx       67    815
Data Rx       0      0
Errors Tx     0
Errors Rx     0
Local name: 2
Remote IP: 198.51.100.10
```

```
Tunnels: 1, Sessions: 1, Switched sessions: 1
State: Enabled
Local IP: 203.0.113.31
Transport: ipUdp, Logical System: default, Router Instance: default
Lockout State: not locked
```

Connections	Totals	Active	Failed
Tunnels	1	1	0
Sessions	1	1	0

	Packets	Bytes
Control Tx	7	462
Control Rx	6	171
Data Tx	0	0
Data Rx	66	798
Errors Tx	0	
Errors Rx	0	

show services l2tp tunnel-switch destination statistics

```
user@host> show services l2tp tunnel-switch destination statistics
Waiting for statistics...
Local name: 2, Tunnels: 1, Sessions: 1
```

	Packets	Bytes
Control Tx	5	452
Control Rx	4	147
Data Tx	0	0
Data Rx	4	54
Errors Tx	0	
Errors Rx	0	

```
Local name: 1, Tunnels: 1, Sessions: 1
```

	Packets	Bytes
Control Tx	4	184
Control Rx	4	243
Data Tx	5	71
Data Rx	0	0
Errors Tx	0	
Errors Rx	0	

show services l2tp tunnel-switch session

Syntax	show services l2tp tunnel-switch session <detail extensive> <statistics>
Release Information	Command introduced in Junos OS Release 13.2.
Description	Display information about L2TP switched tunnel sessions.
Options	<p>none—Display standard information about all active L2TP switched tunnel sessions.</p> <p>detail extensive—(Optional) Display the specified level of output.</p> <p>statistics—(Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with either of the level options, detail or extensive.</p>
Additional Information	
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp tunnel-switch destination on page 763 • show services l2tp tunnel-switch summary on page 772 • show services l2tp tunnel-switch tunnel on page 774
List of Sample Output	show services l2tp tunnel-switch session on page 769 show services l2tp tunnel-switch session detail on page 770 show services l2tp tunnel-switch session extensive on page 770
Output Fields	Table 44 on page 767 lists the output fields for the show services l2tp tunnel-switch session command. Output fields are listed in the approximate order in which they appear.

Table 44: show services l2tp tunnel-switch session Output Fields

Field Name	Field Description	Level of Output
Tunnel local ID	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
Local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	none
Remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	none

Table 44: show services l2tp tunnel-switch session Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	State of the L2TP session: <ul style="list-style-type: none"> • Established—Session is operating. This is the only state supported for the LAC. • closed—Session is being closed. • destroyed—Session is being destroyed. • clean-up—Session is being cleaned up. • lms-ic-accept-new—New session is being accepted. • lms-ic-idle—Session has been created and is idle. • lms-ic-reject-new—New session is being rejected. • lms-ic-wait-connect—Session is waiting for the peer's incoming call connected (ICCN) message. 	All levels
Interface unit	Logical interface for this session.	All levels
Interface Name	(LNS only) Name of an adaptive services interface.	none
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	detail extensive
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	detail extensive
Tunnel switch profile name	Name of a tunnel switch profile.	detail extensive
Mode	(LNS) Mode of the interface representing the session: shared or exclusive . (LAC) Mode of the interface representing the session: shared or dedicated . Only dedicated is currently supported for the LAC.	detail extensive
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	detail extensive
Remote IP	IP address of remote endpoint of the PPP session.	detail extensive
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	detail extensive
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	detail extensive
Bearer type	Type of bearer enabled: <ul style="list-style-type: none"> • 0—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • 1—Digital access requested. • 2—Analog access requested. • 4—Asynchronous Transfer Mode (ATM) bearer support. 	extensive

Table 44: show services l2tp tunnel-switch session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Framing type	Type of framing enabled: <ul style="list-style-type: none"> • 1—Synchronous framing • 2—Asynchronous framing 	extensive
LCP renegotiation	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: On or Off .	extensive
Authentication	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	extensive
Interface ID	(LNS only) Identifier used to look up the logical interface for this session.	extensive
Call serial number	Unique serial number assigned to the call.	extensive
Tx speed	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	extensive
Rx speed	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	extensive
Create time	Day, date, and time when the call was created.	extensive
Up time	Length of time elapsed since the call became active, in hours, minutes, and seconds.	extensive
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	extensive
ToS Reflect	Status of IP ToS value reflection, Disabled or Enabled .	extensive
Statistics since	Date and time when collection of the following statistics began: <ul style="list-style-type: none"> • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. 	extensive

Sample Output

show services l2tp tunnel-switch session

```

user@host> show services l2tp tunnel-switch session
Tunnel local ID: 37602
  Local Remote State           Interface Interface
  ID    ID                    unit      Name
  13545 1      Established        1073741842 si-2/1/0

Tunnel local ID: 37060
  Local Remote State           Interface Interface
  ID    ID                    unit      Name
  58296 1      Established        1073741843 si-2/1/0

```

show services l2tp tunnel-switch session detail

```
user@host> show services l2tp tunnel-switch session detail
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1, Interface unit: 1073741842
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.168.20.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1, Interface unit: 1073741843
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
```

show services l2tp tunnel-switch session extensive

```
user@host> show services l2tp tunnel-switch session extensive
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1
    Interface unit: 1073741842
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.168.20.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Bearer type: 2, Framing type: 1
    LCP renegotiation: On, Authentication: None, Interface ID: si-2/1/0
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:06:50
    Idle time: N/A, ToS Reflect: Disabled
    Statistics since: Fri Jan 18 03:01:11 2013
      Packets      Bytes
      Data Tx      85      1031
      Data Rx       0         0

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1
    Interface unit: 1073741843
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
    Bearer type: 2, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:06:48
    Idle time: N/A
    Statistics since: Fri Jan 18 03:01:14 2013
      Packets      Bytes
```

Data Tx	0	0
Data Rx	84	1014

show services l2tp tunnel-switch summary

Syntax `show services l2tp tunnel-switch summary`
`<statistics>`

Release Information Command introduced in Junos OS Release 13.2.

Description Display L2TP tunnel switch summary information.

Options **none**—Display complete L2TP switched tunnel summary information.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for all switched tunnels and sessions.

Additional Information

Required Privilege Level view

Related Documentation

- [show services l2tp tunnel-switch destination on page 763](#)
- [show services l2tp tunnel-switch session on page 767](#)
- [show services l2tp tunnel-switch tunnel on page 774](#)

List of Sample Output [show services l2tp tunnel-switch summary on page 773](#)

Output Fields [Table 45 on page 772](#) lists the output fields for the **show services l2tp tunnel-switch summary** command. Output fields are listed in the approximate order in which they appear.

Table 45: show services l2tp tunnel-switch summary Output Fields

Field Name	Field Description
Tunnel switch profile name	Name of a tunnel switch profile.
LNS local session id	Identifier assigned by the LNS function on the LTS to the local endpoint of the L2TP session originating on a remote LAC (the first session)
LAC local session id	Identifier assigned by the LAC function on the LTS to the local endpoint of the L2TP session originating on the LTS (the second session).
LNS state	State of the L2TP session (the first session) between a remote LAC and the LNS function on the LTS.
LAC state	State of the L2TP session (the second session) between the LAC function on the LTS and a remote LNS.

Sample Output

show services l2tp tunnel-switch summary

```
user@host> show services l2tp tunnel-switch summary
Tunnel switch profile name: ce-lts-profile
LNS local  LAC local  LNS state    LAC state    Interface
session ID session ID
13545      58296      established  established  si-2/1/0
```

show services l2tp tunnel-switch tunnel

Syntax	show services l2tp tunnel-switch tunnel <detail extensive> <statistics>
Release Information	Command introduced in Junos OS Release 13.2.
Description	Display information about L2TP switched tunnels.
Options	<p>none—Display standard information about all active L2TP tunnels.</p> <p>detail extensive—(Default) Display the specified level of output.</p> <p>statistics—(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with either of the level options, detail or extensive.</p>
Additional Information	
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services l2tp tunnel-switch destination on page 763 • show services l2tp tunnel-switch session on page 767 • show services l2tp tunnel-switch summary on page 772
List of Sample Output	show services l2tp tunnel-switch tunnel on page 777 show services l2tp tunnel-switch tunnel detail on page 777 show services l2tp tunnel-switch tunnel extensive on page 777
Output Fields	Table 46 on page 774 lists the output fields for the show services l2tp tunnel-switch tunnel command. Output fields are listed in the approximate order in which they appear.

Table 46: show services l2tp tunnel-switch tunnel Output Fields

Field Name	Field Description	Level of Output
Local ID	<p>On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS.</p> <p>On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.</p>	none

Table 46: show services l2tp tunnel-switch tunnel Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	none
Remote IP	IP address of the peer endpoint of the tunnel.	All levels
Sessions	Number of L2TP sessions established through the tunnel.	All levels
Switched-sessions	Number of L2TP sessions established by tunnel switching.	All levels
State	State of the L2TP tunnel: <ul style="list-style-type: none"> • cc_responder_accept_new—The tunnel has received and accepted the start control connection request (SCCRQ). • cc_responder_reject_new—The tunnel has received and rejected the SCCRQ. • cc_responder_idle—The tunnel has just been created. • cc_responder_wait_ctl_conn—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • clean-up—The tunnel is being cleaned up. • closed—The tunnel is being closed. • destroyed—The tunnel is being destroyed. • Established—The tunnel is operating. This is the only state supported for the LAC. • Terminate—The tunnel is terminating. • Unknown—The tunnel is not connected to the router. 	All levels
Tunnel local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	detail extensive
Tunnel remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	detail extensive
Tunnel Name	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].	detail extensive
Local IP	IP address of the local endpoint of the tunnel.	detail extensive
Local name	Name used for local tunnel endpoint during tunnel negotiation.	detail extensive
Remote name	Name used for remote tunnel endpoint during tunnel negotiation.	detail extensive

Table 46: show services l2tp tunnel-switch tunnel Output Fields (*continued*)

Field Name	Field Description	Level of Output
Effective Peer Resync Mechanism	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: <ul style="list-style-type: none"> Failover protocol Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol. 	detail extensive
NAS Port Method	(LAC only) Status of interoperation with Cisco LNS devices: <ul style="list-style-type: none"> none—NAS port method is not enabled for interoperation. cisco-avp—NAS port method is enabled for interoperation. 	detail extensive
Tunnel Logical System	Logical system in which the L2TP tunnel is brought up.	detail extensive
Tunnel Routing Instance	Routing instance in which the L2TP tunnel is brought up.	detail extensive
Max sessions	Maximum number of sessions that can be established on this tunnel.	extensive
Window size	Number of control messages that can be sent without receipt of an acknowledgment.	extensive
Hello interval	Interval between the transmission of hello messages, in seconds.	extensive
Create time	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to Unknown and the Create time value resets.	extensive
Up time	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.	extensive
Idle time	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.	extensive
ToS Reflect	Status of IP ToS value reflection, Disabled or Enabled .	extensive
Interface Name	(LNS only) Name of an adaptive services interface.	extensive
Tunnel Group Name	(LNS only) Name of a tunnel group.	extensive
Statistics since	Date and time when collection of the following statistics began: <ul style="list-style-type: none"> Control Tx—Amount of control information transmitted, in packets and bytes. Control Rx—Amount of control information received, in packets and bytes. Data Tx—Amount of data transmitted, in packets and bytes. Data Rx—Amount of data received, in packets and bytes. Errors Tx—Number of errors transmitted, in packets. Errors Rx—Number of errors received, in packets. 	extensive

Sample Output

show services l2tp tunnel-switch tunnel

```
user@host> show services l2tp tunnel-switch tunnel
Local ID Remote ID Remote IP Sessions Switched-sessions State
37602 1 192.168.20.3:1701 1 1 Established
37060 1 198.51.100.10:1701 1 1 Established
```

show services l2tp tunnel-switch tunnel detail

```
user@host> show services l2tp tunnel-switch tunnel detail
Tunnel local ID: 37602, Tunnel remote ID: 1
Remote IP: 192.168.20.3:1701
Sessions: 1, Switched sessions: 1, State: Established
Tunnel Name: 1/1
Local IP: 203.0.113.51:1701
Local name: ce-bras-mx240-f, Remote name: testlac
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: default
Tunnel local ID: 37060, Tunnel remote ID: 1
Remote IP: 198.51.100.10:1701
Sessions: 1, Switched sessions: 1, State: Established
Tunnel Name: 2/1
Local IP: 203.0.113.31:1701
Local name: lns, Remote name: lns
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel-switch tunnel extensive

```
user@host> show services l2tp tunnel-switch tunnel extensive
Waiting for statistics...
Tunnel local ID: 37602, Tunnel remote ID: 1
Remote IP: 192.168.20.3:1701
Sessions: 1, Switched sessions: 1, State: Established
Tunnel Name: 1/1
Local IP: 203.0.113.51:1701
Local name: ce-bras-mx240-f, Remote name: testlac
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: default
Max sessions: 128100, Window size: 4, Hello interval: 60
Create time: Fri Jan 18 03:01:11 2013, Up time: 00:07:49
Idle time: 00:00:00, ToS Reflect: Disabled
Interface Name: si-2/1/0, Tunnel Group Name: ce-l2tp-tunnel-group
Statistics since: Fri Jan 18 03:01:11 2013
      Packets      Bytes
Control Tx         7        259
Control Rx         7        279
Data Tx          97       1175
Data Rx           0           0
Errors Tx          0
Errors Rx          0
Tunnel local ID: 37060, Tunnel remote ID: 1
```

Remote IP: 198.51.100.10:1701
Sessions: 1, Switched sessions: 1, State: Established
Tunnel Name: 2/1
Local IP: 203.0.113.31:1701
Local name: lns, Remote name: lns
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: default
Max sessions: 128100, Window size: 4, Hello interval: 60
Create time: Fri Jan 18 03:01:14 2013, Up time: 00:07:46
Idle time: 00:00:00
Statistics since: Fri Jan 18 03:01:14 2013

	Packets	Bytes
Control Tx	8	482
Control Rx	7	183
Data Tx	0	0
Data Rx	96	1158
Errors Tx	0	
Errors Rx	0	

show services soft-gre tunnel

Syntax	<pre>show services soft-gre tunnel <brief detail extensive> <interface <i>interface-name</i>> <local-ip <i>local-ip-address</i>> <remote-ip <i>remote-ip-address</i>> <statistics> <tunnel-group <i>group-name</i>></pre>
Release Information	Command introduced in Junos OS Release 17.2R1 on MX Series routers.
Description	Display information about dynamic generic routing encapsulation (GRE) tunnels and pseudowire subscriber (psn) interface devices.
Options	<p>none—Display standard information about all active dynamic GRE tunnels.</p> <p>brief detail extensive—(Optional) Display the specified level of detail.</p> <p>interface <i>interface-name</i>—(Optional) Display information for a specific pseudowire subscriber (psn) interface.</p> <p>local-ip <i>local-ip-address</i>—(Optional) Display information for a specific local or source IP address of the GRE tunnel endpoint (Wi-Fi access gateway side).</p> <p>remote-ip <i>remote-ip-address</i>—(Optional) Display information for a specific remote IP address of the GRE tunnel endpoint.</p> <p>statistics—(Optional) Display dynamic GRE tunnel statistics.</p> <p>tunnel-group <i>group-name</i>—(Optional) Display information for a specific dynamic GRE tunnel group.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Wi-Fi Access Gateway Deployment Model Overview on page 344 • Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway on page 345 • Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation on page 349 • soft-gre on page 620
List of Sample Output	<p>show services soft-gre tunnel brief on page 780</p> <p>show services soft-gre tunnel detail on page 780</p> <p>show services soft-gre tunnel extensive on page 781</p>

Output Fields Table 47 on page 780 describes the output fields for the **show services soft-gre tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 47: show services soft-gre tunnel Output Fields

Field Name	Field Description	Level of Output
Interface or Interface Name	Name of the pseudowire subscriber (ps0) interface configured for the GRE tunnel	All levels
Remote IP	Remote IP address of the GRE tunnel endpoint	All levels
Local IP	Local or source IP address of the GRE tunnel endpoint (Wi-Fi access gateway side)	All levels
Subscribers	Number of subscribers accessing the GRE tunnel	All levels
Group Name	Name of the dynamic GRE tunnel group	detail extensive
Routing Instance	Type of routing instance used for the GRE tunnel	detail extensive
Create time	Date and time when the GRE tunnel was created: <i>yyyy-mm-dd hh:mm:ss timezone</i>	detail extensive
Statistics since	Day of the week, date, time, and year when statistics for packets received and transmitted were recorded: <i>Dayofweek Month date hh:mm:ss yyyy</i>	extensive
Statistic	Type of data through the GRE tunnel: Data Rx (data received) and Data Tx (transmitted)	extensive
Packets	Number of data packets received (Data Rx) and transmitted (Data Tx) through the GRE tunnel	extensive
Bytes	Number of data bytes received (Data Rx) and transmitted (Data Tx) through the GRE tunnel	extensive

Sample Output

show services soft-gre tunnel brief

```
user@host> show services soft-gre tunnel brief
Interface      Remote IP      Local IP      Subscribers
ps0.3221225475 50.0.0.10      51.0.0.1      1
```

show services soft-gre tunnel detail

```
user@host> show services soft-gre tunnel detail
Interface Name: ps0.3221225475, Group Name: landslide_v4_1
Local IP: 51.0.0.1
Remote IP: 50.0.0.10
Subscribers: 1
Routing Instance: default
Create time: 2015-03-02 08:06:28 PST
```


show services soft-gre tunnel extensive

```
user@host> show services soft-gre tunnel extensive
Interface Name: ps0.3221225475, Group Name: landslide_v4_1
Local IP: 51.0.0.1
Remote IP: 50.0.0.10
Subscribers: 1
Routing Instance: default
Create time: 2015-03-02 08:06:28 PST
Statistics since: Mon Mar 2 08:06:28 2015
  Statistic      Packets      Bytes
  Data Rx        31         3324
  Data Tx        32         3613
```

show subscribers

Syntax show subscribers
 <detail | extensive | terse>
 <aci-interface-set-name *aci-interface-set-name*>
 <address *address*>
 <agent-circuit-identifier *agent-circuit-identifier-substring*>
 <client-type *client-type*>
 <count>
 <id *session-id* <accounting-statistics>>
 <interface *interface* <accounting-statistics>>
 <logical-system *logical-system*>
 <mac-address *mac-address*>
 <physical-interface *physical-interface-name*>
 <profile-name *profile-name*>
 <routing-instance *routing-instance*>
 <stacked-vlan-id *stacked-vlan-id*>
 <subscriber-state *subscriber-state*>
 <user-name *user-name*>
 <vci *vci-identifier*>
 <vpi *vpi-identifier*>
 <vlan-id *vlan-id*>

Release Information Command introduced in Junos OS Release 9.3.
 Command introduced in Junos OS Release 9.3 for EX Series switches.
 client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
 count option usage with other options introduced in Junos OS Release 10.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
 The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
 Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.
 Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
 Enhanced subscriber management supported in Junos OS Release 15.1R3 for the MX Series.
 accounting-statistics option added in Junos OS Release 15.1R3 and 17.4R1 for MX Series.

Description Display information for active subscribers.

Options **detail | extensive | terse**—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dot1x**—Dot1x clients only.
- **essm**—ESSM clients only.
- **fwauth**—FwAuth (authenticated across a firewall) clients only.
- **l2tp**—L2TP clients only.
- **mlppp**—MLPPP clients only.
- **ppp**—PPP clients only.
- **pppoe**—PPPoE clients only.
- **static**—Static clients only.
- **vlan**—VLAN clients only.
- **vlan-oob**—VLAN out-of-band (ANCP-triggered) clients only.
- **vpls-pw**—VPLS pseudowire clients only.
- **xauth**—Xauth clients only.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id session-id—(Optional) Display a specific subscriber session whose session ID matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

id session-id accounting-statistics—(Optional) Display accurate subscriber accounting statistics for a subscriber session with the specified ID. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface.

interface—(Optional) Display subscribers whose interface matches the specified interface.

interface accounting-statistics—(Optional) Display subscriber accounting statistics for the specified interface. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** **stacked-vlan-id** option to match the outer VLAN tag.



NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level view

- Related Documentation**
- [show subscribers summary on page 812](#)
 - *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
 - *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*
 - *Verifying and Managing Junos OS Enhanced Subscriber Management*

- List of Sample Output**
- [show subscribers \(IPv4\) on page 792](#)
 - [show subscribers \(IPv6\) on page 792](#)
 - [show subscribers \(IPv4 and IPv6 Dual Stack\) on page 792](#)
 - [show subscribers \(Single Session DHCP Dual Stack\) on page 792](#)
 - [show subscribers \(Single Session DHCP Dual Stack detail\) on page 793](#)
 - [show subscribers \(LNS on MX Series Routers\) on page 793](#)
 - [show subscribers \(L2TP Switched Tunnels\) on page 793](#)
 - [show subscribers client-type dhcp detail on page 793](#)
 - [show subscribers client-type dhcp extensive on page 794](#)
 - [show subscribers client-type vlan-oob detail on page 794](#)
 - [show subscribers count on page 795](#)
 - [show subscribers address detail \(IPv6\) on page 795](#)
 - [show subscribers detail \(IPv4\) on page 795](#)
 - [show subscribers detail \(IPv6\) on page 796](#)
 - [show subscribers detail \(pseudowire Interface for GRE Tunnel\) on page 796](#)
 - [show subscribers detail \(IPv6 Static Demux Interface\) on page 796](#)
 - [show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 797](#)
 - [show subscribers detail \(L2TP Switched Tunnels\) on page 797](#)
 - [show subscribers detail \(Tunneled Subscriber\) on page 797](#)
 - [show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 798](#)
 - [show subscribers detail \(ACI Interface Set Session\) on page 798](#)
 - [show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 799](#)
 - [show subscribers extensive on page 799](#)
 - [show subscribers extensive \(Passive Optical Network Circuit Interface Set\) on page 800](#)
 - [show subscribers extensive \(DNS Addresses from Access Profile or Global Configuration\) on page 800](#)
 - [show subscribers extensive \(DNS Addresses from RADIUS\) on page 800](#)
 - [show subscribers extensive \(IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration\) on page 801](#)
 - [show subscribers extensive \(RPF Check Fail Filter\) on page 801](#)
 - [show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 802](#)
 - [show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 802](#)
 - [show subscribers extensive \(ADF Rules \) on page 803](#)
 - [show subscribers extensive \(Effective Shaping-Rate\) on page 803](#)
 - [show subscribers extensive \(PPPoE Subscriber Access Line Rates on page 804](#)
 - [show subscribers extensive \(Subscriber Session Using PCEF Profile\) on page 805](#)
 - [show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 806](#)
 - [show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 806](#)

[show subscribers id accounting-statistics on page 807](#)
[show subscribers interface accounting-statistics on page 807](#)
[show subscribers interface extensive on page 808](#)
[show subscribers logical-system terse on page 809](#)
[show subscribers physical-interface count on page 809](#)
[show subscribers routing-instance inst1 count on page 809](#)
[show subscribers stacked-vlan-id detail on page 809](#)
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 809](#)
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 809](#)
[show subscribers user-name detail on page 809](#)
[show subscribers vlan-id on page 810](#)
[show subscribers vlan-id detail on page 810](#)
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 810](#)
[show subscribers address detail \(Enhanced Subscriber Management\) on page 811](#)

Output Fields [Table 48 on page 786](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 48: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is Tunnel-switched .
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask. (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
Primary DNS Address	IP address of primary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
Secondary DNS Address	IP address of secondary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
IPv6 Primary DNS Address	IPv6 address of primary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
IPv6 Secondary DNS Address	IPv6 address of secondary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
Domain name server inet	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.
Domain name server inet6	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form demux0.nnnn (for example, demux0.3221225472), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is Static or Dynamic .

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
Interface Set	<p>Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line.</p> <p>The name of the interface set uses one of the following prefixes:</p> <ul style="list-style-type: none"> • aci—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • ari—ARI; for example, ari-1033-demux0.3221225524. • aci+ari—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524. • noids—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. <p>NOTE: ACI interface sets are configured with the agent-circuit-identifier autoconfiguration stanza. ALI interface sets are configured with the line-identity autoconfiguration stanza.</p> <p>Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable \$junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.</p>
Interface Set Type	Interface type of the ACI interface set: Dynamic . This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init , Configured , Active , Terminating , Tunneled).
L2TP State	Current state of the L2TP session, Tunneled or Tunnel-switched . When the value is Tunnel-switched , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	<p>For the dhcp client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.</p>
Agent Remote ID	<p>For the dhcp client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.</p>
Accounting Statistics	Actual transmitted subscriber accounting statistics by session ID or interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the actual-transmit-statistics statement is included in the dynamic profile.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 Input Service Set	Input service set in access dynamic profile.

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session ID	ID number for a subscriber service session.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable \$junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).
DSL type	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: ADSL , ADSL2 , ADSL2+ , OTHER , SDSL , VDSL , or VDSL2 .
Frame/Cell Mode	<p>Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091):</p> <ul style="list-style-type: none"> • Cell—When the DSL line type is one of the following: ADSL, ADSL2, or ADSL2+. • Frame—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. <p>The value is stored in the subscriber session database.</p>
Overhead accounting bytes	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
Actual upstream data rate	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).

Table 48: show subscribers Output Fields (*continued*)

Field Name	Field Description
Actual downstream data rate	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
Adjusted downstream data rate	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
Adjusted upstream data rate	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.

Sample Output

show subscribers (IPv4)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/3/0.1073741824  10                  WHOLESALE-CLIENT default:default
demux0.1073741824    203.0.113.10        RETAILER1-CLIENT test1:retailer1
demux0.1073741825    203.0.113.3         RETAILER2-CLIENT test1:retailer1
demux0.1073741826    203.0.113.3         RETAILER2-CLIENT test1:retailer1

```

show subscribers (IPv6)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/0/0.0     2001:db8:c0:0:0:0/74 WHOLESALE-CLIENT default:default
*              2001:db8:1/128      subscriber-25   default:default

```

show subscribers (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
demux0.1073741834 0x8100.1002 0x8100.1 default:default
demux0.1073741835 0x8100.1001 0x8100.1 default:default
pp0.1073741836    203.0.113.13        dualstackuser1@example1.com
default:ASP-1
*                2001:db8:1::/48
*                2001:db8:1:1::/64
pp0.1073741837    203.0.113.33        dualstackuser2@example1.com
default:ASP-1
*                2001:db8:1:2:5::/64

```

show subscribers (Single Session DHCP Dual Stack)

```
user@host> show subscribers
```

```

Interface      IP Address/VLAN ID  User Name      LS:RI
demux0.1073741364 192.168.10.10      dual-stack-retail35 default:default

```

```

2001:db8::100:0:0:0/74
2001:db8:3ffe:0:4::/64
                                default:default

```

show subscribers (Single Session DHCP Dual Stack detail)

```

user@host> show subscribers id 27 detail
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00

```

show subscribers (LNS on MX Series Routers)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-4/0/0.1     192.0.2.0           user@example.com default:default

```

show subscribers (L2TP Switched Tunnels)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-2/1/0.1073741842 Tunnel-switched    user@example.com default:default
si-2/1/0.1073741843 Tunnel-switched    user@example.com default:default

```

show subscribers client-type dhcp detail

```

user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active

```

```

Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 203.0.113.27
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :2560
Login Time: 2009-08-25 14:43:56 PDT

```

show subscribers client-type dhcp extensive

```

user@host> show subscribers client-type dhcp extensive
Type: DHCP
User Name: user
IP Address: 30.11.1.4
IP Netmask: 255.0.0.0
IPv6 Address: 7001:2:3::103
IPv6 Prefix: 7001::/68
Domain name server inet6: abcd::1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:10:94:00:00:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8
00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool

```

show subscribers client-type vlan-oob detail

```

user@host> show subscribers client-type vlan-oob detail
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1

```

```

Logical System: default
Routing Instance: ISP1
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 1234
Session ID: 77
VLAN Id: 126
Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT

```

show subscribers count

```

user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188

```

show subscribers address detail (IPv6)

```

user@host> show subscribers address 203.0.113.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1

```

show subscribers detail (IPv4)

```

user@host> show subscribers detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default

```

```

Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: example :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2

```

show subscribers detail (IPv6)

```

user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (pseudowire Interface for GRE Tunnel)

```

user@host> show subscribers detail

```

Interface	IP Address/VLAN ID	User Name	LS:RI
ps0.3221225484	30.1.0.2		
ps0.3221225485	30.1.0.3		
demux0.3221225486	1		default:default
demux0.3221225487	1		default:default
demux0.3221225488	100.16.0.1		default:default
demux0.3221225489	100.16.0.2		default:default

show subscribers detail (IPv6 Static Demux Interface)

```

user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: user@example.net
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile

```



```

State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT

```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```

user@host> show subscribers detail
Type: L2TP
User Name: user@example.net
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST

```

show subscribers detail (L2TP Switched Tunnels)

```

user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

```

```

Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.31
Remote IP Address: 192.0.2.1
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST

```

show subscribers detail (Tunneled Subscriber)

```

user@host> show subscribers detail

```

```
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST
```

```
Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
```

```
Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
```

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT

```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```

user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.15
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST

```

show subscribers extensive

```

user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48

```

show subscribers extensive (Passive Optical Network Circuit Interface Set)

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic
Interface Set: ot101.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:10:94:00:00:02
State: Active
Radius Accounting ID: jnpr :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|ot101.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: POOL-V4
```

show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (DNS Addresses from RADIUS)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
```

```

IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```

user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

show subscribers extensive (RPF Check Fail Filter)

```

user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof

```

```
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user@example.net
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2016:1:1::1/64
```

```

IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

show subscribers extensive (ADF Rules)

```

user@host> show subscribers extensive
...
Service Session ID: 12
Service Session Name: SERVICE-PROFILE
State: Active
Family: inet
  ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
    Rule 0: 010101000b0101020b020200201811
      from {
        source-address 203.0.113.232;
        destination-address 198.51.100.0/24;
        protocol 17;
      }
      then {
        accept;
      }

```

show subscribers extensive (Effective Shaping-Rate)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST

```

Effective shaping-rate: 31000000k
...

show subscribers extensive (PPPoE Subscriber Access Line Rates)

```

user@host> show subscribers extensive
Type: PPPoE
  IP Address: 192.85.128.1
  IP Netmask: 255.255.255.255
  Logical System: default
  Routing Instance: default
  Interface: pp0.3221225475
  Interface type: Dynamic
  Underlying Interface: demux0.3221225474
  Dynamic Profile Name: pppoe-client-profile-with-cos
  MAC Address: 00:00:6c:07:01:02
  State: Active
  Radius Accounting ID: 4
  Session ID: 4
  PFE Flow ID: 14
  Stacked VLAN Id: 40
  VLAN Id: 1
  Agent Circuit ID: circuit0
  Agent Remote ID: remote0
  Login Time: 2017-04-06 15:52:32 PDT

  User Name: DAVE-L2BSA-SERVICE
  Logical System: default
  Routing Instance: isp-1-subscriber
  Interface: ge-1/2/4.3221225472
  Interface type: Dynamic
  Interface Set: ge-1/2/4
  Underlying Interface: ge-1/2/4
  Core IFL Name: ge-1/3/4.0
  Dynamic Profile Name: L2BSA-88a8-400LL1300V0
  State: Active
  Radius Accounting ID: 1
  Session ID: 1
  PFE Flow ID: 14
  VLAN Id: 13
  VLAN Map Id: 102
  Inner VLAN Map Id: 1
  Agent Circuit ID: circuit-aci-3
  Agent Remote ID: remote49-3
  Login Time: 2017-04-05 16:59:29 EDT
  Service Sessions: 4
  IFL Input Filter Name: L2BSA-CP-400LL1300V0-ge-1/2/4.3221225472-in
  IFL Output Filter Name: L2BSA-CP-400LL1300V0-ge-1/2/4.3221225472-out
  Accounting interval: 900
  DSL type: VDSL
  Frame/Cell Mode: Frame
  Overhead accounting bytes: -10
  Actual upstream data rate: 1024 kbps
  Actual downstream data rate: 4096 kbps
  Adjusted downstream data rate: 3686 kbps
  Adjusted upstream data rate: 922 kbps
  Dynamic configuration:
    junos-vlan-map-id: 102
    Service Session ID: 5
    Service Session Name: SRL-L1
    State: Active

```



```

Family: inet, inet6
IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in
IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out
Service Activation time: 2017-04-05 16:59:30 EDT
Dynamic configuration:
  l2bsa-fwf-in: L2BSA-FWF-in-10048
  l2bsa-fwf-out: L2BSA-FWF-out-25088
  rldown: 25088
  rlup: 10048

```

show subscribers extensive (Subscriber Session Using PCEF Profile)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 5.0.0.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:11:01:00:00:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default

```

```
Service Session ID: 61
Service Session Name: pcef-serv-prof
State: Active
Family: inet
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: limit-fb
Service Activation time: 2017-03-28 08:31:19 PDT
Dynamic configuration:
  pcef-prof: pcef-prof-1
  pcef-rule1: limit-fb
  svc-filt: svc-filt-1
  svc-set: tdf-service-set
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address:
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
```

```

State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:52
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers id accounting-statistics

```

user@host> show subscribers id 601 accounting-statistics
Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

show subscribers interface accounting-statistics

```

user@host> show subscribers interface pp0.3221226949 accounting-statistics
Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322

```

```
Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: user@test.example.com
IP Address: 192.0.2.0
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface      IP Address/VLAN ID  User Name      LS:RI
demux0.1073741825  203.0.113.3      RETAILER1-CLIENT test1:retailer1
demux0.1073741826  203.0.113.4      RETAILER2-CLIENT test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
```

```
IP Address: 203.0.113.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface          IP Address          User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
```

```
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers address detail (Enhanced Subscriber Management)

```
user@host> show subscribers address 203.0.113.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show subscribers summary

Syntax show subscribers summary
 <all>
 <detail | extensive | terse>
 <count>
 <physical-interface *physical-interface-name*>
 <logical-system *logical-system* pic | port | routing-instance *routing-instance* | slot>

Release Information Command introduced in Junos OS Release 10.2.

Description Display summary information for subscribers.

Options **none**—Display summary information by state and client type for all subscribers.

all—(Optional) Display summary information by state, client type, and LS:RI.

detail | extensive | terse—(Not supported on MX Series routers) (Optional) Display the specified level of output.

count—(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system *logical-system*—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface *physical-interface-name*—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type, and LS:RI.

pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.

port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.

routing-instance *routing-instance*—(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.



NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level view

Related Documentation • [show subscribers on page 782](#)

List of Sample Output [show subscribers summary on page 814](#)
[show subscribers summary all on page 815](#)
[show subscribers summary physical-interface on page 815](#)
[show subscribers summary physical-interface pic on page 815](#)
[show subscribers summary physical-interface port on page 816](#)
[show subscribers summary physical-interface slot on page 816](#)
[show subscribers summary pic on page 816](#)
[show subscribers summary pic \(Aggregated Ethernet Interfaces\) on page 816](#)
[show subscribers summary port on page 816](#)
[show subscribers summary port \(Pseudowire Interfaces\) on page 817](#)
[show subscribers summary port extensive on page 817](#)
[show subscribers summary slot on page 817](#)
[show subscribers summary terse on page 817](#)

Output Fields [Table 49 on page 813](#) lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 49: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
Subscribers by State	Number of subscribers summarized by state. The summary information includes the following: <ul style="list-style-type: none"> Init—Number of subscriber currently in the initialization state. Configured—Number of configured subscribers. Active—Number of active subscribers. Terminating—Number of subscribers currently terminating. Terminated—Number of terminated subscribers. Total—Total number of subscribers for all states. 	detail none
Subscribers by Client Type	Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPOE, STATIC-INTERFACE, VLAN, and VLAN-OOB. Also displays the total number of subscribers for all client types (Total).	detail extensive none
Subscribers by LS:RI	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).	detail none
Subscribers by Connection Type	Number of subscribers summarized by connection type, Cross-connected or Terminated .	extensive

Table 49: show subscribers summary Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface	<p>Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.</p> <p>The * character indicates a continuation of addresses for the same session.</p> <p>For aggregated Ethernet interfaces, the output of the summary (pic port slot) options prefixes the interface name with ae0:.</p> <p>For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and redundant logical tunnel (RLT) anchor interface. For example:</p> <p>ps0: 1t-2/1/0 ps1: r1t0: 1t-4/0/0</p>	All levels
Count	<p>Count of subscribers displayed for each PIC, port, or slot when those options are specified with the summary option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.</p> <p>Multiple pseudowire interfaces can share a given logical tunnel or redundant logical tunnel anchor interface. Starting in Junos OS Release 18.1R1, the field displays subscribers per individual pseudowire interface.</p> <p>In earlier releases, the field displays the same number of subscribers for all pseudowire interfaces that share the same tunnel interface as their anchor point.</p>	detail extensive none
Total Subscribers	Total number of subscribers for all physical interfaces, all PICs, all ports, or all LS:RI slots.	detail extensive none
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>	terse
User Name	Name of subscriber.	terse
LS:RI	Logical system and routing instance associated with the subscriber.	terse

Sample Output

show subscribers summary

```
user@host> show subscribers summary
```

```
Subscribers by State
Init          3
Configured    2
Active       183
Terminating   2
Terminated    1

TOTAL        191
```

```

Subscribers by Client Type
DHCP      107
PPP       76
VLAN       8
VLAN-OOB   2
TOTAL     193

```

show subscribers summary all

```

user@host> show subscribers summary all
Subscribers by State
Init      3
Configured 2
Active    183
Terminating 2
Terminated 1

TOTAL     191

Subscribers by Client Type
DHCP      107
PPP       76
VLAN       8

TOTAL     191

Subscribers by LS:RI
default:default 1
default:ri1     28
default:ri2     16
ls1:default     22
ls1:riA         38
ls1:riB         44
logsysX:routinstY 42

TOTAL     191

```

show subscribers summary physical-interface

```

user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
Active: 3998
Total: 3998

Subscribers by Client Type
DHCP: 3998
Total: 3998

Subscribers by LS:RI
default:default: 3998
Total: 3998

```

show subscribers summary physical-interface pic

```

user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
Active: 4825
Total: 4825

```

```
Subscribers by Client Type
DHCP: 4825
Total: 4825
```

```
Subscribers by LS:RI
default:default: 4825
Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
Active: 4825
Total: 4825

Subscribers by Client Type
DHCP: 4825
Total: 4825

Subscribers by LS:RI
default:default: 4825
Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
Active: 4825
Total: 4825

Subscribers by Client Type
DHCP: 4825
Total: 4825

Subscribers by LS:RI
default:default: 4825
Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
Interface      Count
ge-1/0         1000
ge-1/3         1000

Total Subscribers: 2000
```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
Interface      Count
ae0: ge-1/0    801
ae0: ge-1/3    801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
```

Interface	Count
ge-5/0/1	201
ge-5/0/2	301

Total Subscribers: 502

show subscribers summary port (Pseudowire Interfaces)

```
user@host> show subscribers summary port
ps0: lt-2/1/0 10
ps1: lt-2/1/0 20
```

Total Subscribers: 30

show subscribers summary port extensive

```
user@host>show subscribers summary port extensive
Interface: ge-5/0/1
Count: 201
Detail:
Subscribers by Client Type
  DHCP: 100
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 200
  Cross-connected: 1
```

```
Interface: ge-5/0/2
Count: 301
Detail:
Subscribers by Client Type
  DHCP: 200
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 300
  Cross-connected: 1
```

Total Subscribers: 502

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface      Count
ge-1           2000
```

Total Subscribers: 2000

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/3/0.1073741824  100                WHOLESALER-CLIENT default:default
demux0.1073741824    203.0.113.10        RETAILER1-CLIENT default:default
demux0.1073741825    203.0.113.13        RETAILER1-CLIENT test1:retailer1
demux0.1073741826    203.0.113.213       RETAILER2-CLIENT test1:retailer2
```

show system subscriber-management statistics

Syntax	show system subscriber-management statistics <all> <dhcp> <pppoe>
Release Information	Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management. Enhanced I/O Statistics introduced as part of Extensive output in Junos OS Release 15.1R4 on MX Series routers for enhanced subscriber management.
Description	Display statistics for the specified option. You can customize the output by including one or more optional filters in the command. With the exception of the extensive option, all filter options can be combined in a single command.
Options	all —(Optional) Display packet statistics for all protocol.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Dropped Packets and Untransmitted Traffic Using show Commands</i>
List of Sample Output	show system subscriber-management statistics all on page 819 show system subscriber-management statistics pppoe on page 820 show system subscriber-management statistics dhcp on page 820 show system subscriber-management statistics dhcp extensive on page 821 show system subscriber-management statistics extensive on page 823
Output Fields	Table 50 on page 818 lists the output fields for the show system subscriber-management statistics command. Output fields are listed in the approximate order in which they appear.

Table 50: show system subscriber-management statistics Output Fields

Field Name	Field Description
Rx Statistics	Statistics for packets received.
Tx Statistics	Statistics for packets sent.
Enhanced I/O Statistics	Statistics for visibility into packet drops from the queue.
Error Statistics	Includes connection packets, flow control, and messages and packets sent to and received from the daemon.

Table 50: show system subscriber-management statistics Output Fields (*continued*)

Field Name	Field Description
ERA discards	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
padis	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
padrs	PPPoE Active Discovery Request packets.
ppp	Point-to-Point Protocol packets.
router solicitations	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.
router advertisements	Number of router advertisements sent or received.
route solicit response packet	Number of router solicitation responses sent or received.

Sample Output

The following examples displays packet statistics accumulated for DHCP and PPPoE since the last time the session manager was cleared.

show system subscriber-management statistics all

```

user@host> show system subscriber-management statistics all
  user@host> show system subscriber-management statistics all
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
-----
                        Packet Statistics
-----
I/O Statistics:
-----
    Rx Statistics
      packets                : 784711
    Tx Statistics
      packets                : 7013122
  Layer 3 Statistics
    Rx Statistics
      packets                : 356218
    Tx Statistics
      packets                : 6604660

DHCP Statistics:
-----
    Rx Statistics
      packets                : 320008
      ERA discards           : 6274
    Tx Statistics
      transmit request packets : 320482

```

```
sent packets                : 320482
Error Statistics
Connection Statistics
no connection packets       : 0
```

PPPoE Statistics:

```
-----
Rx Statistics
  packets                : 486165
  padis                  : 36768
  padrs                  : 35421
  ppp packets            : 341787
  ERA discards           : 8249
Tx Statistics
  packets                : 70842
  send failures          : 6240
```

show system subscriber-management statistics pppoe

```
user@host> show system subscriber-management statistics pppoe
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
```

Packet Statistics

I/O Statistics:

```
-----
Rx Statistics
  packets                : 784711
Tx Statistics
  packets                : 7013122
Layer 3 Statistics
Rx Statistics
  packets                : 356218
Tx Statistics
  packets                : 6604660
```

PPPoE Statistics:

```
-----
Rx Statistics
  packets                : 486165
  padis                  : 36768
  padrs                  : 35421
  ppp packets            : 341787
  ERA discards           : 8249
Tx Statistics
  packets                : 70842
  send failures          : 6240
```

show system subscriber-management statistics dhcp

```
user@host> show system subscriber-management statistics dhcp
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
```

Packet Statistics

I/O Statistics:

```
-----
Rx Statistics
```



```

        packets                : 784711
    Tx Statistics
        packets                : 7013122
Layer 3 Statistics
    Rx Statistics
        packets                : 356218
    Tx Statistics
        packets                : 6604660

```

DHCP Statistics:

```

-----
Rx Statistics
    packets                : 320008
    ERA discards           : 6274
Tx Statistics
    transmit request packets : 320482
    sent packets           : 320482
Error Statistics
Connection Statistics
    no connection packets   : 0

```

show system subscriber-management statistics dhcp extensive

```

user@host> show system subscriber-management statistics dhcp extensive
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015

```

Packet Statistics

I/O Statistics:

```

-----
Rx Statistics
    packets                : 784711
Tx Statistics
    packets                : 7013122
Buffer Statistics
    allocations            : 7032618
    frees                  : 7032624
    allocation failures     : 0
Layer 3 Statistics
    Rx Statistics
        packets                : 356218
    Tx Statistics
        packets                : 6604660
PFE Event Statistics
    packets                : 0

```

DHCP Statistics:

```

-----
Rx Statistics
    packets                : 320008
    ERA discards           : 6274
Tx Statistics
    transmit request packets : 320482
    sent packets           : 320482
DHCPv4 Rx Statistics
    total packets          : 0
DHCPv4 Tx Statistics
    total packets          : 0
DHCPv6 Rx Statistics

```

total packets	: 320008
solicit	: 36250
request	: 36382
renew	: 247376
ERA discards	: 6274
DHCPv6 Tx Statistics	
total packets	: 320482
advertise	: 36382
reply	: 284100
Error Statistics	
Connection Statistics	
no connection packets	: 0
connection down events	: 0
connection up events	: 0
flow control invoked	: 0
flow control released	: 0
packets sent to daemon	: 320008
packets received from daemon	: 320482
messages sent to daemon	: 0
messages received from daemon	: 0
notifies while not connected	: 0

NET Statistics:

ICMP6 Statistics

Rx Statistics

packets:	: 36271
router solicitations	: 36271

Tx Statistics

packets:	: 6284178
router advertisements	: 6284178
route solicit response packet	: 36271

Management Statistics:

dvlan	: 33912
dvlan adds	: 33912
pppoe	: 143651
pppoe add	: 35750
pppoe changes	: 107901
ip flow	: 143633
ip flow add	: 107883

Management Config Status:

gres state enabled state	: 1
shmlog disabled state	: 0
Rx Statistics	
packets	: 167361
ERA discards	: 15116
Tx Statistics	
transmit request packets	: 150903
sent packets	: 150903
DHCPv4 Rx Statistics	
total packets	: 167361
discover	: 91910
request	: 75451
ERA discards	: 15116

show system subscriber-management statistics extensive

```
user@host> show system subscriber-management statistics extensive
```

```
Session Manager started @ Tue Nov 3 10:00:57 2015
```

```
Session Manager cleared @ Tue Nov 3 11:10:01 2015
```

```
-----
                        Packet Statistics
-----
```

```
I/O Statistics:
-----
```

```
Rx Statistics
```

```
  packets : 784711
```

```
Tx Statistics
```

```
  packets : 7013122
```

```
Buffer Statistics
```

```
  allocations : 7032618
```

```
  frees : 7032624
```

```
  allocation failures : 0
```

```
Layer 3 Statistics
```

```
Rx Statistics
```

```
  packets : 356218
```

```
Tx Statistics
```

```
  packets : 6604660
```

```
PFE Event Statistics
```

```
  packets : 0
```

```
-----
Enhanced I/O Statistics:
-----
```

```
bbe_io_rcv l2 : 0
```

```
bbe_io_rcv l3 : 0
```

```
bbe_io_rcv l3 v4 : 0
```

```
io low queue drops :12
```

```
io mlow queue drops :0
```

```
io medium queue drops :0
```

```
io high queue drops :0
```

show system subscriber-management summary

Syntax	show system subscriber-management summary
Release Information	Command introduced in Junos OS Release 11.1. Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Description	Display complete subscriber management database summary information.
Options	none —This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>show database-replication statistics</i> • <i>show database-replication summary</i>
List of Sample Output	show system subscriber-management summary on page 825 show system subscriber-management summary (Enhanced Subscriber Management) on page 826
Output Fields	Table 51 on page 824 lists the output fields for the show system subscriber-management summary command. Output fields are listed in the approximate order in which they appear.

Table 51: show system subscriber-management summary Output Fields

Field Name	Field Description
Graceful Restart	State of graceful Routing Engine switchover (GRES): <ul style="list-style-type: none"> • Enabled • Disabled (Enhanced subscriber management for MX Series routers) The name of this field is Graceful Switchover .
Mastership	State of the Routing Engine: <ul style="list-style-type: none"> • Master • Standby
Database	State of the subscriber management database: <ul style="list-style-type: none"> • Available • Init • Not-available

Table 51: show system subscriber-management summary Output Fields (continued)

Field Name	Field Description
Standby	<p>(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine:</p> <ul style="list-style-type: none"> • Connected—Connected but not synchronized • Disconnected—Not connected • Resync (nn%)—Connected and <i>nn</i> percent synchronized with the master Routing Engine • Synchronized—Synchronized with the master Routing Engine
Chassisd ISSU State	<p>State of unified ISSU chassis daemon:</p> <ul style="list-style-type: none"> • ABORT • DAEMON_ISSU_PREPARE • DAEMON_ISSU_PREPARE_DONE • DAEMON_SWITCHOVER_PREPARE • DAEMON_SWITCHOVER_PREPARE_DONE • FRU_ISSU • FRU_ISSU_DONE • IDLE • UNKNOWN
ISSU State	<p>State of unified ISSU aggregate daemon:</p> <ul style="list-style-type: none"> • ABORT • IDLE • PREPARE • READY • SWITCHOVER_PREPARE • SWITCHOVER_READY • UNKNOWN
ISSU Wait	<p>Amount of time, in seconds, requested by a daemon to perform cleanup. If multiple daemons request time, the displayed value is the highest wait time requested by a daemon.</p>

Sample Output

show system subscriber-management summary

```

user@host> show system subscriber-management summary
General:
  Graceful Restart      Enabled
  Mastership            Master
  Database              Available
  Chassisd ISSU State   DAEMON_ISSU_PREPARE
  ISSU State            PREPARE
  ISSU Wait             198

```

show system subscriber-management summary (Enhanced Subscriber Management)

```
user@host> show system subscriber-management summary
General:
  Graceful Switchover      Enabled
  Mastership               Master
  Database                 Available
  Standby                  Resync (75%)
  Chassisd ISSU State      IDLE
  ISSU State               IDLE
  ISSU Wait                0
```

test services l2tp tunnel

Syntax	test services l2tp tunnel user <i>user-name</i> <password <i>user-password</i>> <tunnel-name <i>name</i>>
Release Information	Command introduced in Junos OS Release 11.4.
Description	(MX Series routers only) Test and verify Layer 2 Tunneling Protocol (L2TP) tunnel configurations from the L2TP access concentrator (LAC). The test determines whether the user can be authenticated and tunneled according to the L2TP configuration. The establishment of all tunnels associated with the user is tested. You can optionally specify a particular tunnel to test for the user.
Options	<p>user <i>user-name</i>—Name of the user under test. You must use an existing configured username, although it can be created solely for testing a tunnel configuration.</p> <p>password <i>user-password</i>—(Optional) Authentication password for the specified user. If you omit this option, the test generates a dummy password—<i>testpass</i>—for the user.</p> <p>tunnel-name <i>name</i>—(Optional) Name of a tunnel to test.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Testing L2TP Tunnel Configurations from the LAC on page 300
List of Sample Output	<p>test services l2tp tunnel user (User authentication fails) on page 828</p> <p>test services l2tp tunnel user (Multiple tunnels tested) on page 828</p> <p>test services l2tp tunnel user tunnel-name (Specific tunnel tested) on page 828</p>
Output Fields	Table 52 on page 827 lists the output fields for the test services l2tp tunnel command. Output fields are listed in the approximate order in which they appear.

Table 52: test services l2tp tunnel Output Fields

Field Name	Field Description
Tunnel-name	Name of the tunnel as configured in the local tunnel profile.
Tunnel-peer	IP address of the tunnel's remote peer (the L2TP network server [LNS]).
Logical-System	Logical system in which the tunnel is created.
Routing-Instance	Routing instance in which the tunnel is created.

Table 52: test services l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
Status	Status of the tunnel.

Sample Output

test services l2tp tunnel user (User authentication fails)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication failed
```

test services l2tp tunnel user (Multiple tunnels tested)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication success, l2tp tunneled
Tunnel-name  Tunnel-peer  Logical-System  Routing-Instance  Status
test1tunnel  192.168.2.3  default        default          Up
test2tunnel  198.51.100.243  default        default          Peer unresponsive
test3tunnel  198.51.100.251  default        test            Up
```

test services l2tp tunnel user tunnel-name (Specific tunnel tested)

```
user@host> test services l2tp tunnel user testuser@example.com tunnel-name test1tunnel
Subscriber: testuser@example.com, authentication success, l2tp tunneled
Tunnel-name  Tunnel-peer  Logical-System  Routing-Instance  Status
test1tunnel  192.168.2.3  default        default          Up
```