JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Broadband Subscriber Access Protocols User Guide

Published
2020-06-19

JUNOS

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. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Junos® OS Broadband Subscriber Access Protocols User Guide
Copyright © 2020 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://support.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**

Using the Examples in This Manual | **xxiv**

- Merging a Full Example | **xxiv**
- Merging a Snippet | **xxv**

Documentation Conventions | **xxv**

Documentation Feedback | **xxviii**

Requesting Technical Support | **xxviii**

- Self-Help Online Tools and Resources | **xxix**
- Creating a Service Request with JTAC | **xxix**

1 **Broadband Subscriber Access Network Overview**

Broadband Subscriber Access Network Overview | 31

- Subscriber Access Network Overview | **31**
- Multiservice Access Node Overview | **32**
- Ethernet MSAN Aggregation Options | **33**
  - Direct Connection | **34**
  - Ethernet Aggregation Switch Connection | **34**
  - Ring Aggregation Connection | **35**
- LDP Pseudowire Autosensing Overview | **35**
  - Pseudowire Ingress Termination Background | **36**
  - Pseudowire Autosensing Approach | **37**
  - Sample Configuration | **38**
- Layer 2 Services on Pseudowire Service Interface Overview | **38**
  - Traffic from Customer LAN to MPLS | **39**
  - Traffic from Service Edge to Customer LAN | **40**
  - Pseudowire Service Interfaces | **41**
  - Sample Configuration | **41**
- Broadband Access Service Delivery Options | **47**
  - Digital Subscriber Line | **47**
  - Active Ethernet | **48**

Passive Optical Networking | 48

Hybrid Fiber Coaxial | 49

Broadband Delivery and FTTx | 49

Understanding BNG Support for Cascading DSLAM Deployments Over Bonded DSL Channels | 50

Benefits of Cascading DSLAM Deployments Over Bonded DSL Channels | 50

4-Level Scheduler Hierarchy | 50

Use Cases of Cascading DSLAM Deployments Over Bonded DSL Channels | 51

Bonded DSL for Copper-To-The-Building (CuTTB) | 52

Hybrid PON + G.fast | 53

Supported Features | 53

Detection of Backhaul Line Identifiers and Autogeneration of Intermediate Node Interface Sets | 54

High Availability for Subscriber Access Networks | 58

Unified ISSU for High Availability in Subscriber Access Networks | 58

Verifying and Monitoring Subscriber Management Unified ISSU State | 59

Graceful Routing Engine Switchover for Subscriber Access Networks | 60

DHCP | 60

L2TP | 61

Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover | 62

Routes for DHCP and PPP Subscriber Access Networks | 63

Access and Access-Internal Routes for Subscriber Management | 63

Configuring Dynamic Access Routes for Subscriber Management | 64

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers | 66

Suppressing DHCP Access, Access-Internal, and Destination Routes | 68

Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default | 68

Verifying the Configuration of Access and Access-Internal Routes for DHCP and PPP Subscribers | 69

Subscribers with Identical Framed Routes | 70

2

DHCP Subscriber Access Networks

DHCP Subscriber Access Networks Overview | 73

DHCP and Subscriber Management Overview | **73**

Extended DHCP Local Server and Subscriber Management Overview | **74**

Extended DHCP Relay and Subscriber Management Overview | **74**

DHCP Relay Proxy and Subscriber Management Overview | **74**

Subscriber Access Operation Flow Using DHCP Relay | **74**

Defining Various Levels of Services for DHCP Subscribers | **75**

Example: Configuring a Tiered Service Profile for Subscriber Access | **76**

DHCP Snooping for Network Security | 81

DHCP Snooping Support | **81**

What is DHCP Snooping | **81**

Benefits of DHCP Snooping | **81**

Configuring DHCP Snooping | **82**

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server | **83**

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent | **84**

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent | **91**

Disabling DHCP Snooping Filters | **93**

Example: Configuring DHCP Snooping Support for DHCP Relay Agent | **95**

Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent | **97**

Preventing DHCP Spoofing | **103**

DHCPv4 Duplicate Client Management | 104

DHCPv4 Duplicate Client In Subnet Overview | **105**

Guidelines for Configuring Support for DHCPv4 Duplicate Clients | **105**

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information | **106**

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces | **108**

DHCPv6 Duplicate Client Management | 110

DHCPv6 Duplicate Client DUIDs | **110**

Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs | **110**

3

PPP Subscriber Access Networks

PPP Subscriber Access Networks Overview | 114

Dynamic Profiles for PPP Subscriber Interfaces Overview | **114**

Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests | **115**

Benefits of PPP Fast Keepalives | **115**

How PPP Fast Keepalive Processing Works | **116**

Statistics Display for PPP Fast Keepalive | **116**

Effect of Changing the Forwarding Class Configuration | **117**

Ignoring a Magic Number Mismatch | **117**

RADIUS-Sourced Connection Status Updates to CPE Devices | **118**

Message and Option Formats | **120**

Configuring Dynamic Profiles for PPP | **123**

Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges | **124**

How to Configure RADIUS-Sourced Connection Status Updates to CPE Devices | **125**

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | **126**

Migrating Static PPP Subscriber Configurations to Dynamic Profiles Overview | **127**

Local Authentication | **127**

CPE-Sourced Address Assignment | **127**

Tag2 Route Attribute | **128**

Benefits | **128**

Configuring Local Authentication in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | **128**

Configuring Tag2 Attributes in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | **130**

Configuring Dynamic Authentication for PPP Subscribers | **131**

Modifying the CHAP Challenge Length | **133**

Example: Minimum PPPoE Dynamic Profile | **134**

Verifying and Managing PPP Configuration for Subscriber Management | **135**

PPP Network Control Protocol Negotiation | 136

PPP Network Control Protocol Negotiation Mode Overview | **136**

PPP NCP Negotiation Modes | **137**

PPP NCP Negotiation Mode Supported Configurations | **138**

PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers | 138

PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers | 138

PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations | 139

Controlling the Negotiation Order of PPP Authentication Protocols | 140

Configuring the PPP Network Control Protocol Negotiation Mode | 142

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses | 144

Tracing PPP Service Events for Troubleshooting | 145

Configuring the PPP Service Trace Log Filename | 147

Configuring the Number and Size of PPP Service Log Files | 147

Configuring Access to the PPP Service Log File | 148

Configuring a Regular Expression for PPP Service Messages to Be Logged | 148

Configuring Subscriber Filtering for PPP Service Trace Operations | 149

Configuring the PPP Service Tracing Flags | 150

Configuring the Severity Level to Filter Which PPP Service Messages Are Logged | 150

4 **L2TP Subscriber Access Networks**

L2TP for Subscriber Access Overview | 153

L2TP for Subscriber Access Overview | 153

L2TP Terminology | 156

L2TP Implementation | 157

Sequence of Events on the LAC | 157

Sequence of Events on the LNS | 158

Retransmission of L2TP Control Messages | **159**

Configuring Retransmission Attributes for L2TP Control Messages | **160**

Enabling Tunnel and Global Counters for SNMP Statistics Collection | **163**

Verifying and Managing L2TP for Subscriber Access | **164**

L2TP Tunnel Switching For Multiple-Domain Networks | 166

L2TP Tunnel Switching Overview | **166**

Application of Tunnel Switch Profiles | **168**

Termination of Tunnel-Switched Sessions on the LTS | **168**

Tunnel Switching Actions for L2TP AVPs at the Switching Boundary | **170**

Configuring L2TP Tunnel Switching | **175**

Setting the L2TP Receive Window Size | **177**

Setting the L2TP Tunnel Idle Timeout | **178**

Setting the L2TP Destruct Timeout | **178**

Configuring the L2TP Destination Lockout Timeout | **180**

Removing an L2TP Destination from the Destination Lockout List | **181**

Configuring L2TP Drain | **181**

Using the Same L2TP Tunnel for Injection and Duplication of IP Packets | **182**

L2TP LAC Subscriber Configuration | 183

Configuring an L2TP LAC | **183**

Configuring How the LAC Responds to Address and Port Changes Requested by the LNS | **186**

LAC Interoperation with Third-Party LNS Devices | **189**

Globally Configuring the LAC to Interoperate with Cisco LNS Devices | **189**

L2TP LAC Tunneling for Subscribers | 191

LAC Tunnel Selection Overview | **191**

Selection When Failover Between Preference Levels Is Configured | **194**

Selection When Failover Within a Preference Level Is Configured | **199**

Selection When Distributing the Session Load Across Multiple LNSs | **201**

L2TP Session Limits Overview | **207**

Scenario 1: Chassis Limit | **208**

Scenario 2: Tunnel Limit | **209**

Scenario 3: Tunnel Group Limit | **209**

Scenario 4: Session-Limit Group Limit | **210**

Scenario 5: Individual Client Limit | 212

Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | 213

Setting the Format for the Tunnel Name | 215

Configuring a Tunnel Profile for Subscriber Access | 216

Configuring the L2TP LAC Tunnel Selection Parameters | 219

Configuring LAC Tunnel Selection Failover Within a Preference Level | 220

Configuring Weighted Load Balancing for LAC Tunnel Sessions | 221

Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions | 221

Enabling the LAC for IPv6 Services | 222

Testing L2TP Tunnel Configurations from the LAC | 222

L2TP Subscriber Access Lines and Connection Speeds | 225

Subscriber Access Line Information Handling by the LAC and LNS Overview | 225

Access Line Information Forwarding | 226

Access Line Information AVPs | 227

Connection Speed Updates on the LAC | 234

Connection Speed Updates on the LNS | 235

Interaction Between Global and Per-Destination Configurations | 236

Transmission of Tx and Rx Connection Speeds from LAC to LNS | 236

Methods for Determining the Speed Values Reported to the LNS | 237

Determining Initial Connect Speeds | 241

Fallback Mechanism for Connect Speed Values | 241

Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal | 244

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS | 244

Configuring the Reporting and Processing of Subscriber Access Line Information | 247

Preventing the LAC from Sending Calling Number AVP 22 to the LNS | 252

Override the Calling-Station-ID Format for the Calling Number AVP | 253

Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds | 255

L2TP LNS Inline Service Interfaces | 260

Configuring an L2TP LNS with Inline Service Interfaces | 261

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

Configuring an L2TP Access Profile on the LNS | 268

Configuring a AAA Local Access Profile on the LNS | 270

Configuring an Address-Assignment Pool for L2TP LNS with Inline Services | 271

Configuring the L2TP LNS Peer Interface | 273

Enabling Inline Service Interfaces | 273

Configuring an Inline Service Interface for L2TP LNS | 275

Configuring Options for the LNS Inline Services Logical Interface | 276

LNS 1:1 Stateful Redundancy Overview | 277

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278

Verifying LNS Aggregated Inline Service Interface 1:1 Redundancy | 280

L2TP Session Limits and Load Balancing for Service Interfaces | 283

Session Limits on Service Interfaces | 283

Session Load Balancing Across Service Interfaces | 284

Example: Configuring an L2TP LNS | 287

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301

Applying Services to an L2TP Session Without Using RADIUS | 304

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions | 313

Configuring a Dynamic Profile for Dynamic LNS Sessions | 314

IP Packet Reassembly on Inline Service Interfaces | 317

IP Packet Fragment Reassembly for L2TP Overview | 317

Configuring IP Inline Reassembly for L2TP | 318

Peer Resynchronization After an L2TP Failover | 320

L2TP Failover and Peer Resynchronization | 320

Configuring the L2TP Peer Resynchronization Method | 321

Tracing L2TP Events for Troubleshooting | 323

Configuring the L2TP Trace Log Filename | 324

Configuring the Number and Size of L2TP Log Files | 325

Configuring Access to the L2TP Log File | 325

Configuring a Regular Expression for L2TP Messages to Be Logged | 326

Configuring Subscriber Filtering for L2TP Trace Operations | 326

Configuring the L2TP Tracing Flags | 328

Configuring the Severity Level to Filter Which L2TP Messages Are Logged | 328

5 Configuring MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces | 331

Pseudowire Subscriber Logical Interfaces Overview | 331

Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview | 334

Configuring a Pseudowire Subscriber Logical Interface | 338

Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router | 340

Configuring a Pseudowire Subscriber Logical Interface Device | 341

Changing the Anchor Point for a Pseudowire Subscriber Logical Interface Device | 344

Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface | 346

Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces | 347

Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces | 348

Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface | 350

6 Wi-Fi Access Gateways

Wi-Fi Access Gateways | 357

Wi-Fi Access Gateway Overview | 358

Wi-Fi Access Gateway Deployment Model Overview | 360

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway | 361

Dynamic VLAN-Tagged Subscribers | 362

Untagged Subscribers | 362

Wi-Fi Access Gateway Configuration Overview | 362

Configuring a Pseudowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway | 363

Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation | 365

Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways | 367

Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways | 371

7 Fixed Wireless Access Networks

Fixed Wireless Access Networks | 374

Fixed Wireless Access Network Overview | 374

References | 378

3GPP Fixed Wireless Access Terminology | 378

Benefits of Fixed Wireless Access | 383

How to Configure Fixed Wireless Access | 383

Verifying and Monitoring Fixed Wireless Access | 387

Tracing Fixed Wireless Access Events for Troubleshooting | 388

Configuring the Fixed Wireless Access Trace Log Filename | 389

Configuring the Number and Size of Fixed Wireless Access Log Files | 389

Configuring Access to the Fixed Wireless Access Log File | 390

Configuring a Regular Expression for Fixed Wireless Access Messages to Be Logged | 391

Configuring the Fixed Wireless Access Tracing Flags | 391

8

Configuration Statements

aaa-access-profile (L2TP LNS) | 401

aaa-context (AAA Options) | 402

aaa-options (Access Profile) | 403

aaa-options (PPP Profile) | 404

access (Dynamic Access Routes) | 406

access-internal (Dynamic Access-Internal Routes) | 408

access-line (Access-Line Rate Adjustment) | 410

access-line-information (L2TP) | 426

access-profile (AAA Options) | 428

address (L2TP Destination) | 429

address (L2TP Tunnel Destination) | 430

address (LNS Local Gateway) | 431

address (Tunnel Profile Remote Gateway) | 432

address (Tunnel Profile Source Gateway) | 433

address-change-immediate-update | 434

aggregated-inline-services-options (Aggregated Inline Services) | 435

allow-snooped-clients | 437

always-write-option-82 | 439

anchor-point (Pseudowire Subscriber Interfaces) | 441

assignment-id-format (L2TP LAC) | 444

authentication (Static and Dynamic PPP) | 445

avp (L2TP Tunnel Switching) | 446

bandwidth (Inline Services) | 447

bandwidth (Tunnel Services) | 449

bearer-type (L2TP Tunnel Switching) | 451

bfd | 452

calling-number (L2TP Tunnel Switching) | 454

challenge-length (Static and Dynamic PPP) | 455

chap | 457

chap (Dynamic PPP) | 459

chap (L2TP) | 460

cisco-nas-port-info (L2TP Tunnel Switching) | 461

client | 462

delimiter (Access Profile) | 466

destination (L2TP) | 467

destination-equal-load-balancing (L2TP LAC) | 468

destruct-timeout (L2TP) | 469

detection-time | 470

device-count (Pseudowire Subscriber Interfaces) | 471

dhcp-local-server | 472

dhcp-relay | 489

dhcpv6 (DHCP Local Server) | 508

dhcpv6 (DHCP Relay Agent) | 517

dial-options | 526

dial-options (Dynamic Profiles) | 528

disable-calling-number-avp (L2TP LAC) | 529

disable-failover-protocol (L2TP) | 530

drain | 532

dual-stack-group (DHCP Local Server) | 533

dual-stack-group (DHCP Relay Agent) | 535

duplicate-clients (DHCPv6 Local Server and Relay Agent) | 538

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent) | 540

dynamic-profile (L2TP) | 542

dynamic-profile (PPP) | 543

dynamic-profiles | 544

enable-ipv6-services-for-lac (L2TP) | 558

enable-snmp-tunnel-statistics (L2TP) | 559

enforce-strict-scale-limit-license (Subscriber Management) | 560

equals (Dynamic Profile) | 561

failover-resync | 562

failover-within-preference (L2TP LAC) | 563

failure-action | 564

flexible-vlan-tagging | 566

forward-snooped-clients (DHCP Local Server) | 567

forward-snooped-clients (DHCP Relay Agent) | 568

fpc (MX Series 5G Universal Routing Platforms) | 570

gateway-name (LNS Local Gateway) | 572

gateway-name (Tunnel Profile Remote Gateway) | 573

gateway-name (Tunnel Profile Source Gateway) | 574

gres-route-flush-delay (Subscriber Management) | 575

group (DHCP Local Server) | 576

group (DHCP Relay Agent) | 582

group-profile (Group Profile) | 588

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | 590

holddown-interval | 592

hello-interval (L2TP) | 593

identification (Tunnel Profile) | 594

idle-timeout (Access) | 595

idle-timeout (L2TP) | 597

ignore-magic-number-mismatch (Access Group Profile) | 598

ignore-magic-number-mismatch (Dynamic Profiles) | 600

initiate-ncp (Dynamic and Static PPP) | 602

inline-services (PIC level) | 604

input-hierarchical-policer | 605

interface (Dynamic Routing Instances) | 606

interface (Service Interfaces) | 607

interface-id | 608

interfaces (Static and Dynamic Subscribers) | 609

ip-reassembly | 616

ip-reassembly (L2TP) | 617

ip-reassembly-rules (Service Set) | 618

ipcp-suggest-dns-option | 619

keepalive | 620

keepalives | 621

keepalives (Dynamic Profiles) | 623

l2tp | 625

l2tp (Profile) | 629

l2tp-access-profile | 633

l2tp-maximum-session (Service Interfaces) | 634

layer2-liveness-detection (Receive) | 636

layer2-liveness-detection (Send) | 638

lcp-renegotiation | 640

liveness-detection | 641

local-authentication (Dynamic PPP Options) | 643

local-gateway (L2TP LNS) | 644

lockout-timeout (L2TP Destination Lockout) | 645

logical-system (Tunnel Profile) | 646

mac | 647

mac-address (Dynamic Access-Internal Routes) | 648

match-direction (IP Reassembly Rule) | 649

maximum-sessions (L2TP) | 650

maximum-sessions-per-tunnel | 651

max-sessions (Tunnel Profile) | 652

medium (Tunnel Profile) | 653

method | 654

metric (Dynamic Access-Internal Routes) | 656

minimum-interval | 657

minimum-receive-interval | 659

minimum-retransmission-timeout (L2TP Tunnel) | 660

mtu | 661

multiplier | 665

name (L2TP Destination) | 666

name (L2TP Tunnel Destination) | 667

no-adaptation | 668

nas-port-method (L2TP LAC) | 669

nas-port-method (Tunnel Profile) | 670

next-hop (Dynamic Access Routes) | 671

next-hop-service | 673

no-allow-snooped-clients | 675

no-gratuitous-arp-request | 676

no-snoop (DHCP Local Server and Relay Agent) | 677

on-demand-ip-address | 679

options (Access Profile) | 681

override (RADIUS Options) | 692

overrides (DHCP Relay Agent) | 694

overrides (Enhanced Subscriber Management) | 696

pap | 700

pap (Dynamic PPP) | 701

pap (L2TP) | 702

parse-direction (Access Profile) | 703

pic (M Series and T Series Routers) | 705

pool (Service Interfaces) | 707

pp0 (Dynamic PPPoE) | 708

ppp (Group Profile) | 710

ppp-options | 713

ppp-options (Dynamic PPP) | 715

ppp-options (L2TP) | 718

preference (Subscriber Management) | 720

preference (Tunnel Profile) | 721

primary-interface (Aggregated Inline Services) | 722

profile (Access) | 723

proxy-mode | 730

ps0 (Pseudowire Subscriber Interfaces) | 731

pseudowire-service (Pseudowire Subscriber Interfaces) | 732

qualified-next-hop (Dynamic Access-Internal Routes) | 733

radius (Access Profile) | 734

reject-unauthorized-ipv6cp | 738

relay-option-82 | 739

remote-gateway (Tunnel Profile) | 741

report-ingress-shaping-rate (Dynamic CoS Interfaces) | 742

request services l2tp destination unlock | 744

retransmission-count-established (L2TP) | 745

retransmission-count-not-established (L2TP) | 746

route (Access) | 747

route (Access Internal) | 749

route-suppression (DHCP Local Server and Relay Agent) | 751

routing-instance (Tunnel Profile) | 752

routing-instance (L2TP Destination) | 753

routing-instance (L2TP Tunnel Destination) | 754

routing-instances (Dynamic Profiles) | 755

routing-options (Dynamic Profiles) | 757

rule (IP Reassembly) | 759

rx-connect-speed-when-equal (L2TP LAC) | 760

rx-window-size (L2TP) | 761

secondary-interface (Aggregated Inline Services) | 762

secret (Tunnel Profile) | 763

service-device-pool (L2TP) | 764

service-device-pools (Service Interfaces) | 765

service-interface (L2TP Processing) | 766

service-profile (L2TP) | 767

service-rate-limiter (Access) | 769

session-mode | 771

session-options | 772

sessions-limit-group (L2TP) | 776

soft-gre | 777

source-gateway (Tunnel Profile) | 779

stacked-vlan-tagging | 780

statistics (Access Profile) | 781

strip-user-name (Access Profile) | 782

subscriber-context (AAA Options) | 783

subscriber-management (Subscriber Management) | 784

tag (Access) | 787

tag2 (Dynamic Access Routes) | 788

threshold (detection-time) | 789

threshold (transmit-interval) | 791

tos-reflect (L2TP) | 792

trace (DHCP Relay Agent) | 793

traceoptions (Services L2TP) | 795

traceoptions (Protocols PPP Service) | 800

traceoptions (Subscriber Management) | 804

transmit-interval | 806

tunnel (L2TP) | 807

tunnel (Tunnel Profile) | 809

tunnel-group | 811

tunnel-profile (L2TP Tunnel Switching) | 813

tunnel-profile (Tunnel Profile) | 814

tunnel-switch-profile (L2TP Tunnel Switching, Application) | 815

tunnel-switch-profile (L2TP Tunnel Switching, Definition) | 816

tx-address-change (L2TP LAC) | 817

tx-connect-speed-method (L2TP LAC) | 819

type (Tunnel Profile) | 822

unit (Dynamic PPPoE) | 823

unit (Dynamic Profiles Standard Interface) | 826

untagged | 830

username-include (Local Authentication) | 831

version (BFD) | 833

weighted-load-balancing (L2TP LAC) | 834

vlan-id (Dynamic Profiles) | 835

vlan-tagging | 836

vlan-tagging (Dynamic) | 839

vlan-tags | 840

9

Operational Commands

clear services l2tp destination | 844

clear services l2tp destination lockout | 846

clear services l2tp session | 848

clear services l2tp session statistics | 851

clear services l2tp tunnel | 853

clear services l2tp tunnel statistics | 856

request interface (revert | switchover) (Aggregated Inline Service Interfaces) | 858

show ancp subscriber | 860

show bfd subscriber session | 870

show dynamic-profile session | 876

show interfaces ps0 (Pseudowire Subscriber Interfaces) | 882

show interfaces redundancy | 888

show ppp interface | 892

show ppp statistics | 912

show ppp summary | 921

show services inline ip-reassembly statistics | 923

show services l2tp client | 929

show services l2tp destination | 931

show services l2tp destination lockout | 936

show services l2tp session | 938

show services l2tp session-limit-group | 949

show services l2tp summary | 951

show services l2tp tunnel | 959

show services l2tp tunnel-group | 967

show services l2tp tunnel-switch destination | 969

show services l2tp tunnel-switch session | 974

show services l2tp tunnel-switch summary | 980

show services l2tp tunnel-switch tunnel | 982

show services soft-gre tunnel | 988

show subscribers | 992

show subscribers summary | 1040

show system subscriber-management statistics | 1049

show system subscriber-management summary | 1059

test services l2tp tunnel | 1063

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | **xxiii**
- Using the Examples in This Manual | **xxiv**
- Documentation Conventions | **xxv**
- Documentation Feedback | **xxviii**
- Requesting Technical Support | **xxviii**

Use this guide to understand how to configure the primary methods for accessing the subscriber network:

- DHCP provides IP address configuration and service provisioning.
- PPP enables a point-to-point direct connection to the network and service provider. Dynamic profiles apply configurations and services to authenticated subscribers.
- L2TP separates the termination of access technologies from the termination of PPP and subsequent access to a network. This separation enables service providers to outsource their access technologies. L2TP provides ISPs the capability to supply VPN service; private enterprises can reduce or avoid investment in access technologies for remote workers.
- MPLS pseudowire interfaces extend MPLS domains from the access-aggregation network to the service edge.
- Wi-Fi access gateways provide public Wi-Fi access from residential or business Wi-Fi networks so that mobile subscribers can be authenticated and connected regardless of their physical location.

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.

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.xsl;
        }
    }
}
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.xsl; }
```

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 xxvi 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`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*

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

Convention	Description	Examples
Italic text like this	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** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  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.	• 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 so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

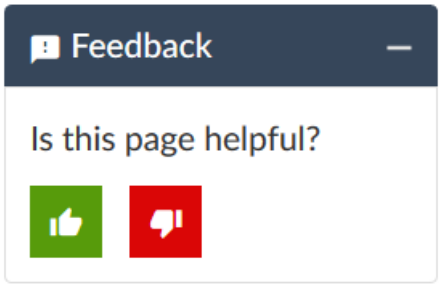

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- 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 Juniper Care or Partner Support Services 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/
- Create a service request online: https://myjuniper.juniper.net

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- 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://support.juniper.net/support/requesting-support/.

1

CHAPTER

Broadband Subscriber Access Network Overview

Broadband Subscriber Access Network Overview | **31**

High Availability for Subscriber Access Networks | **58**

Routes for DHCP and PPP Subscriber Access Networks | **63**

Subscribers with Identical Framed Routes | **70**

Broadband Subscriber Access Network Overview

IN THIS SECTION

- Subscriber Access Network Overview | 31
- Multiservice Access Node Overview | 32
- Ethernet MSAN Aggregation Options | 33
- LDP Pseudowire Autosensing Overview | 35
- Layer 2 Services on Pseudowire Service Interface Overview | 38
- Broadband Access Service Delivery Options | 47
- Broadband Delivery and FTTx | 49
- Understanding BNG Support for Cascading DSLAM Deployments Over Bonded DSL Channels | 50
- Detection of Backhaul Line Identifiers and Autogeneration of Intermediate Node Interface Sets | 54

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 32 shows an example of a basic subscriber access network.

Figure 1: Subscriber Access Network Example

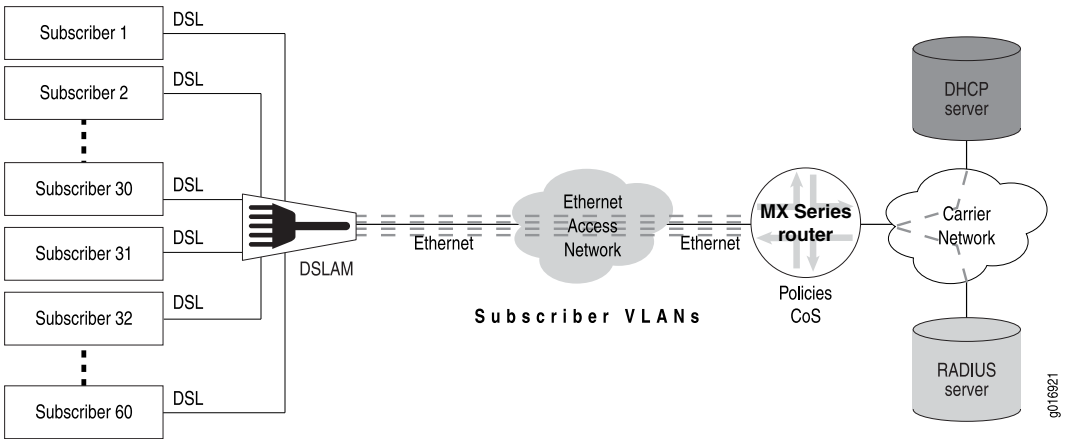

NOTE: This feature requires a license. To understand more about Subscriber Access Licensing, see, Subscriber Access Licensing Overview. Please refer to the Juniper Licensing Guide for general information about License Management. Please refer to the product Data Sheets at MX Series Routers for details, or contact your Juniper Account Team or Juniper Partner.

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 33:

Figure 2: Choosing an MSAN Type

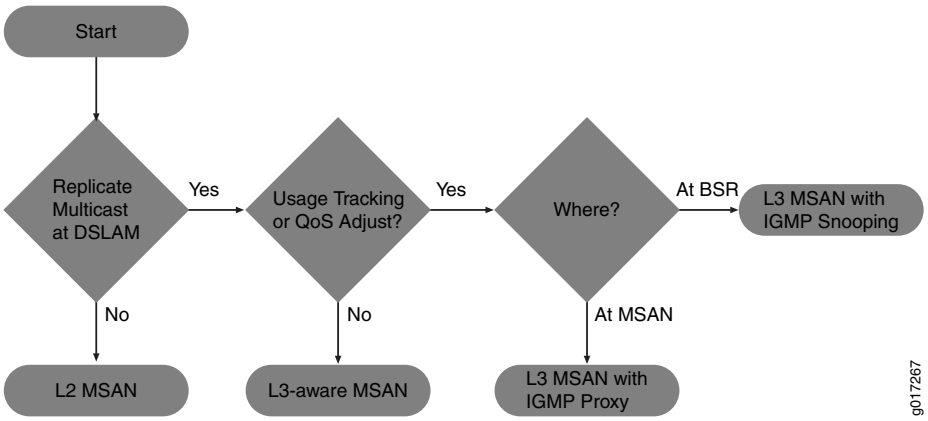

Ethernet MSAN Aggregation Options

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 33 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 *Ethernet Networking User 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 *Ethernet Networking User Guide for MX Series Routers*.

LDP Pseudowire Autosensing Overview

IN THIS SECTION

- Pseudowire Ingress Termination Background | **36**
- Pseudowire Autosensing Approach | **37**
- Sample Configuration | **38**

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 signals 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

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 37

Figure 3: Basic Control Flow of Pseudowire Autosensing

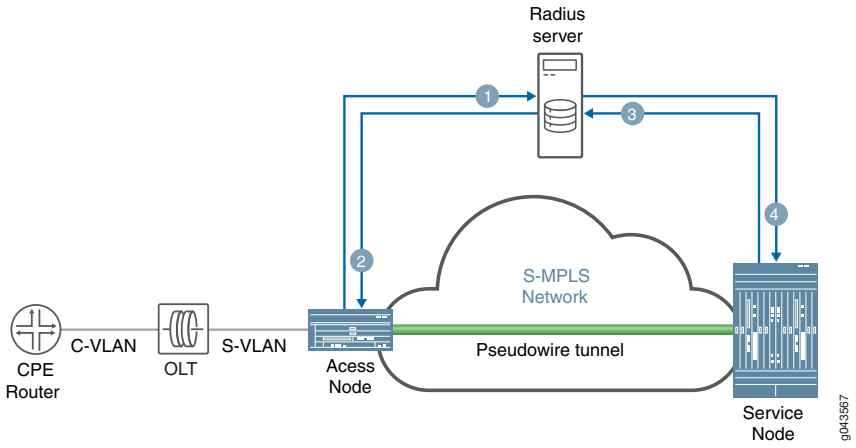

The basic control flow procedure of pseudowire autosensing is as follows:

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 0

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

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 MLPS access between an Ethernet aggregation device and a service edge device with the pseudowire service on transport logical interface psn.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. Starting in Junos OS Release 17.1R1, the pseudowire logical tunnel interfaces support Ethernet VPLS, Ethernet bridge, VLAN VPLS, and VLAN bridge encapsulation next hops to exit Layer 2 traffic. Starting in Junos OS Release 18.4R1, the Layer 2 service support with the pseudowire service logical interfaces is extended to pseudowire service interfaces anchored over redundant logical tunnel interfaces as well. These Layer 2 services are supported only on pseudowire service on service logical interfaces (psn.1 to psn.n) and not on transport logical interface (psn.0). The Layer 2 output features such as VLAN manipulations and others are enabled on the pseudowire service interfaces. The traffic sent out of the 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 (EAD 1) 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 40, when the traffic comes from EAD 1 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

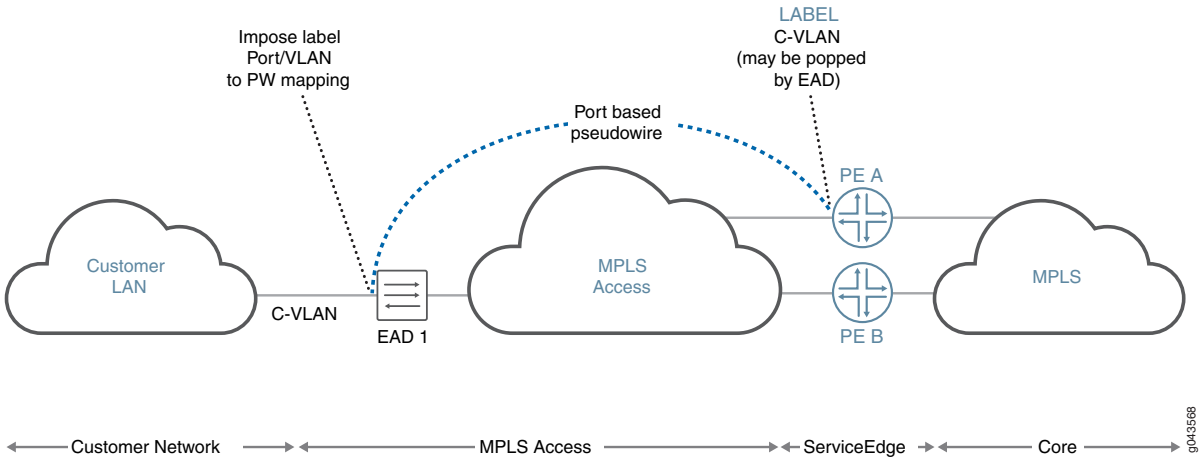

When traffic enters pseudowire service on the service logical interface ps0.n, where n>0, the following steps are performed.

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 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;
        }
    }
    l2-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 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;
        }
    }
    l2-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 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;
      }
    }
    l2-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;
            }
         }
      }
   }
```

Broadband Access Service Delivery Options

IN THIS SECTION

- Digital Subscriber Line | **47**
- Active Ethernet | **48**
- Passive Optical Networking | **48**
- Hybrid Fiber Coaxial | **49**

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

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:

- ATM PON (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—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 (GEPON) 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.

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.

Understanding BNG Support for Cascading DSLAM Deployments Over Bonded DSL Channels

Junos OS supports configuring and maintaining the access lines between access nodes and their ANCP subscribers using DSL access multiplexer as the broadband access technology for Copper-to-the-Building (CuTTB) and Fiber-to-the-Building (FTTB). When multiple subscribers share the same access line, the access line could be one of the following types:

- PON, Fiber-to-the-Building (FTTB)
- Bonded DSL Copper-To-The-Building (CTTB)

Starting in Junos OS Release 18.2R1, Passive Optical Network (PON) access technologies are supported with four levels of quality-of-service (QoS) scheduler hierarchy for residential subscribers in a BBE deployment. This feature extends the Access Node Control Protocol (ANCP) implementation to handle network configuration for residential customers that use PON as the broadband access technology for both CuTTB and FTTB. ANCP uses a statically controlled traffic-control profile on the interface-set for shaping at the subscriber level at the intermediate node to which the subscribers are connected. New DSL types are provided to support access line rate adjustment for the new access technologies.

A new RADIUS VSA, **Inner-Tag-Protocol-Id** 26-211 is introduced to fetch the inner VLAN Tag Protocol Identifier value for L2BSA subscribers to enable maintaining one dynamic profile instead of two separate dynamic profiles. A new Junos OS dynamic profile variable *$junos-inner-vlan-tag-protocol-id* allows a VLAN map's **inner-tag-protocol-id** to be set by RADIUS or a predefined default value provided in the configuration.

Benefits of Cascading DSLAM Deployments Over Bonded DSL Channels

This feature is useful to support access network deployments where multiple subscribers share the same access line aggegrated by an intermediate node between the access node and the home routing gateways. Another benefit is to conserve Layer 2 CoS nodes. Typically a dummy Layer 2 node is created for each residential household, which could exhaust Layer 2 CoS resources. Therefore, network models using bonded DSL, G.Fast, and PON access models can conserve Layer 2 CoS nodes.

4-Level Scheduler Hierarchy

Junos OS supports 4-Level QoS scheduler hierarchy minimally supporting residential and L2BSA access over Copper-to-the-Building (CTTB) or Fiber-to-the-Building access network deployments. The following QoS scheduler hierarchy levels are supported:

- Level 1 Port (Physical interface or AE)
- Level 2 Access Line (Logical interface set, represents a collection of subscribers sharing a given access line aggregated by an intermediate node)

- Level 3 Subscriber sessions
- Level 4 Queues (services)

Figure 5: Scheduler Hierarchy

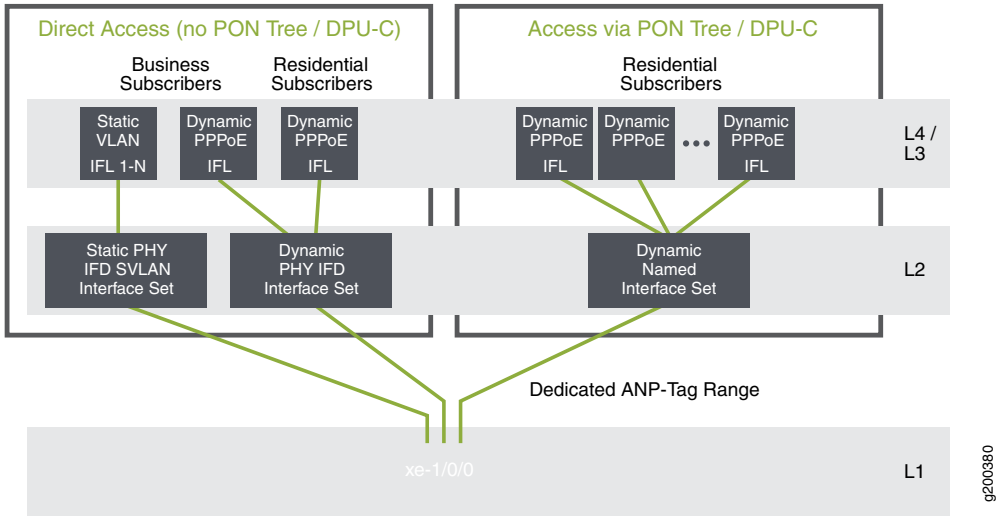

In Figure 5 on page 51, residential and L2BSA access require only 4-level scheduler hierarchy. Business subscriber access is currently not supported and hence 4-level scheduler hierarchy is sufficient for CuTTB and PON services targeted to an apartment building.

Use Cases of Cascading DSLAM Deployments Over Bonded DSL Channels

Bonded DSL for copper to the building (CuTTB) introduces an intermediate node Distribution Point Unit-Copper (DPU-C) between the DSL access multiplexer (DSLAM) and a cluster of subscribers at the customer location. Shared access line deployment models may be of type Passive-Optical-Network (PON) or bonded DSL copper lines. Example intermediate nodes are listed below:

- DPU-C - bonded DSL for Copper-To-The-Building (CTTB)
- ONU - PON (Fiber-to-the-Building (FTTB)
- Hybrid PON and G.Fast

Bonded DSL for Copper-To-The-Building (CuTTB)

Figure 6: Bonded DSL/CuTTB

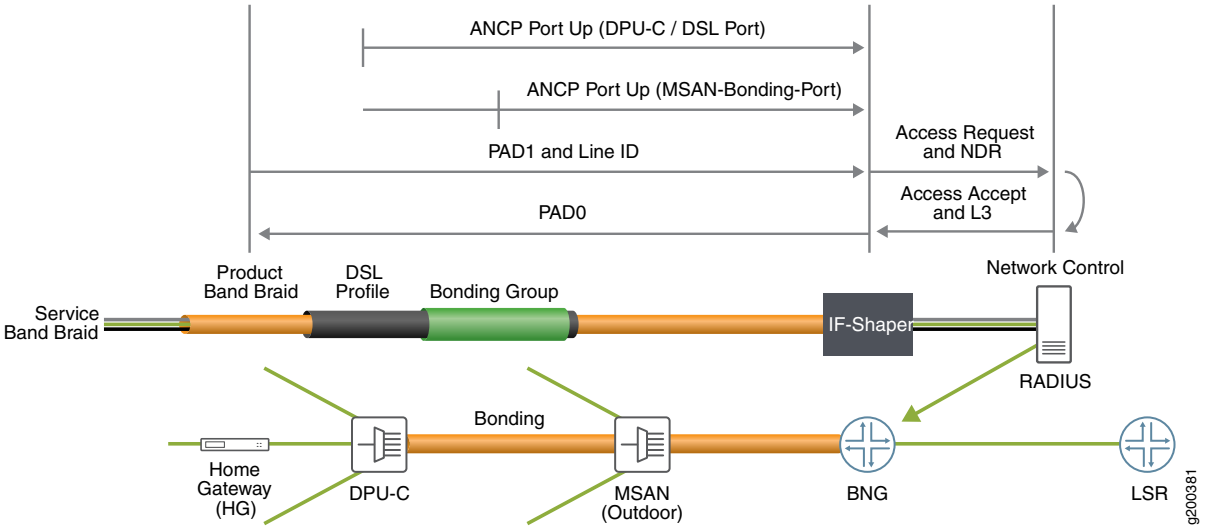

In Figure 6 on page 52, each DPU-C has an ANCP session to report access line parameters of individual subscribers connected to the node. The MSAN also has an ANCP session to report access line parameters of the bonded DSL access line to the DPU-C. All subscribers connected to the DPU-C are thus subject to the DSL access-line downstream rate, the DPU-C subscribers are grouped together in an interface set. You can adjust the speeds reported in this Port-Up and apply to the CoS node for the corresponding interface ste maintaining the semantics of the CoS adjustment control profile that is used for individual subscriber lines. The access model consists of a hybrid of bonded DSL access and conventional unbonded access. The DPU-C and the Multi Service Access Node (MSAN) ANCP sessions are completely independent and the PPPoE-IA tags only reflect the attributes reported in the dPU-C ANCP session

Hybrid PON + G.fast

Figure 7: Hybrid PON + G.fast

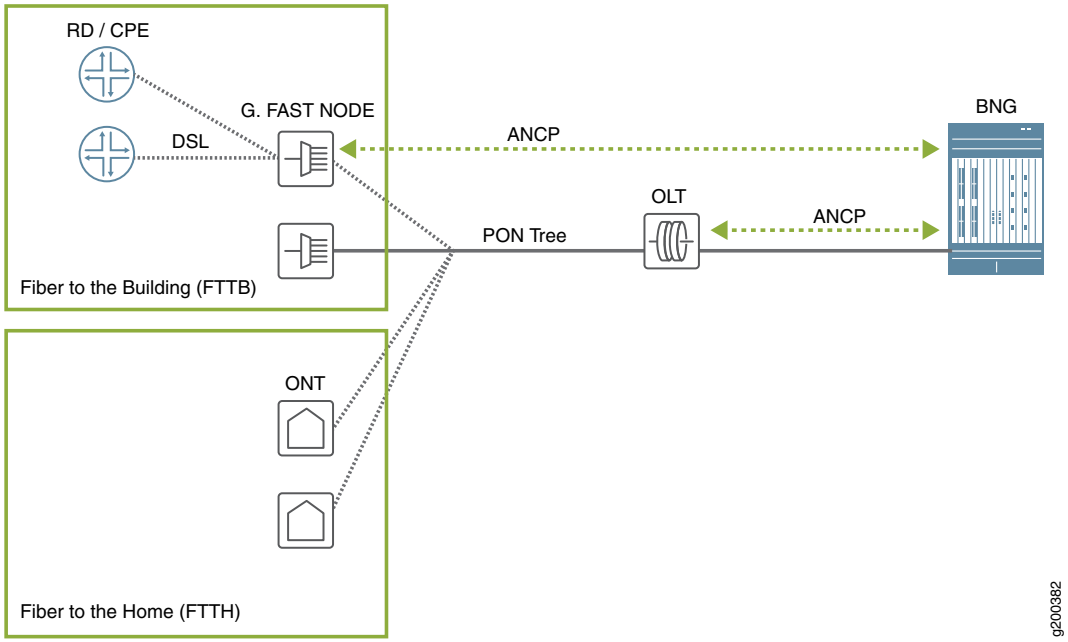

In Figure 7 on page 53, the OLT has an ANCP session with the BNG and proxies for all downstream native PON nodes. G.fast DSL subscribers are connected to an intermediate node, which has a PON connection to the intermediate ONU in front of the OLT.

A hybrid access network connects DSL based subscriber lines using both PON access and G.fast nodes with an intermediate node between the OLT and home gateways (HGs). Both businesses and residences are connected to the intermediate node, which is the PON leaf. Shaping is required both at the subscriber level and at the PON leaf level. Ths G.fast subscribers are associated with the intermediate ONU like a native PON subscriber. New DSL type TLVs are supported by the AN and their values are reported in the ANCP Port-Up for the correstpnding subscriber access line. However, it is still not possible to distinguish between an intermediate node and a conventional connection for a given PPPoE session.

Supported Features

- Support ANCP-based traffic shaping on dynamic iflsets.
- Preservation of PPP0E-IA and ANCP independence by CLI configuration for residential subscribers.
- New Juniper VSA, ERX-Inner-Vlan-Tag-Protocol-Id (4874-26-211) is supported to source the inner VLAN Tag Protocol Identifier value for L2BSA subscribers as an optimization to maintain two, separate

dynamic profiles, one for TPID - 0x88a8 and one for 0x8100, and sourcing the desired value by returning 26-4874-174 (Client-rofile-Name) in the Access-Accept.

- The following additional type values for the DSL type TLV are supported. All subscribers include these DSL type TLVs in the PPPoE PADR messages's PPPoE IA tags.
  - (8) G.fast
  - (9) VDSL2 Annex Q
  - (10) SDSL bonded
  - (11) VDSL2 bonded
  - (12) G,fast bonded
  - (13) VDSL2 Annex Q bonded

Detection of Backhaul Line Identifiers and Autogeneration of Intermediate Node Interface Sets

Starting in Junos OS Release 18.4R1, you can configure the router to detect a logical intermediate node in an access network. The node identifies subscribers that are connected to the same shared media, such as a PON tree or a bonded copper line that connects to a DPU-C for CuTTB. When you configure this detection, the router parses the ANCP Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x03) that is received either in the ANCP Port Up message or PPPoE PADR IA tags. If the TLV string begins with the **#** character, the string is a backhaul line identifier that is unique across the network to identify the bonded DSL line or the PON tree. The same string is reported in the TLV or IA for all subscribers connected to that DPU-C or PON.

The portion of the string after the **#** character represents the logical intermediate node. It is used as the name of the dynamic interface set for the CoS Level 2 node that groups the subscribers using that intermediate node. This interface set is known as the parent interface set. Every PPPoE or VLAN (L2BSA) logical interface with the same value for TLV 0x03 is a member of that interface set.

> **NOTE:** The TLV value must match the requirements for interface set naming; it can include alphanumeric characters and the following special characters:
>
> # % / = + - : ; @ . _

This portion of the string also sets the value of the $junos-aggregation-interface-set-name predefined variable in the dynamic profile. This value is used as the name of a CoS Level 2 interface set that groups

the subscribers sharing that string. It overrides the predefined variable default, which uses the value of $junos-phy-ifd-interface-set-name as the name of the interface set.

For example, if the value of the TLV string is #TEST-DPU-C-100, the value of the predefined variable—and consequently the name of the interface set—becomes TEST-DPU-C-100.

> **NOTE:** The Access-Loop-Remote-ID (TLV (0x02) is similarly parsed for the **#** character, but the resulting string is not used in the current release.

> **NOTE:** Intermediate node detection is supported only for 4-level scheduler hierarchies, so business access is limited to conventional DSL access MPCs.

Before you begin, you must confirm that your existing access nodes or IAs are not already inserting strings that begin with the **#** character. Because this is a system-level configuration, parsing applies to all ANCP access nodes and PPPoE IAs globally. The leading **#** character is not configurable. Parsing is disabled by default in case some providers use that character for some other purpose.

To enable parsing of the Access-Aggregation-Circuit-ID-ASCII TLV and setting the interface set name:

1. Specify detection of hierarchical access networks and extraction of the node string.

   ```
   [edit system access-line]
   user@host# set hierarchical-access-network-detection
   ```

2. Configure the dynamic profile to use the Access-Aggregation-Circuit-ID-ASCII string for the interface set name.

   ```
   [edit dynamic-profile interfaces]
   user@host# set interface-set $junos-aggregation-interface-set-name
   ```

The following sample configuration shows a dynamic profile for L2BSA subscribers. Three things to note here are the following:

- A default value of $junos-phy-ifd-interface-set-name is defined for the $junos-aggregation-interface-set-name predefined variable.
- The name of the interface set is configured to be the value of $junos-aggregation-interface-set-name.
- The CoS scheduler configuration specifies an interface named with the value of $junos-aggregation-interface-set-name.

When **hierarchical-access-network-detection** is configured for the access lines, then the name of the Level 2 scheduler interface set is determined as follows:

- When TLV 0x03 begins with **#**, then $junos-aggregation-interface-set-name is the remainder of the string, excluding the initial **#**.
- When TLV 0x03 begins with any other character, then $junos-aggregation-interface-set-name is the value of $junos-phy-ifd-interface-set-name.

```
[edit dynamic-profiles L2BSA-subscriber]
predefined-variable-defaults {
   aggregation-interface-set-name phy-ifd-interface-set-name;
   cos-shaping-rate 1g;
   cos-scheduler-map schedmap_L2BSA;
   inner-vlan-tag-protocol-id 0x88a8;
}
routing-instances {
   "$junos-routing-instance" {
      interface "$junos-interface-name";
   }
}
interfaces {
   interface-set $junos-aggregation-interface-set-name {
      interface "$junos-interface-ifd-name" {
         unit "$junos-interface-unit";
      }
   }
   "$junos-interface-ifd-name" {
      unit "$junos-interface-unit" {
         encapsulation vlan-vpls;
         no-traps;
         vlan-id "$junos-vlan-id";
         input-vlan-map {
            swap-push;
            inner-tag-protocol-id "$junos-inner-vlan-tag-protocol-id"
            vlan-id "$junos-vlan-map-id";
            inner-vlan-id "$junos-inner-vlan-map-id";
         }
         output-vlan-map {
            pop-swap;
            inner-tag-protocol-id 0x8100;
         }
         family vpls;
      }
   }
```

```
}
class-of-service {
    traffic-control-profiles {
        L2BSAShaper {
            scheduler-map "$junos-cos-scheduler-map";
            shaping-rate "$junos-cos-shaping-rate" burst-size 17k;
            overhead-accounting frame-mode cell-mode-bytes 6;
        }
        L2iflsetShaper {
            shaping-rate 1G burst-size 17k;
        }
    }
    interfaces {
        "$junos-interface-ifd-name" {
            unit "$junos-interface-unit" {
                output-traffic-control-profile L2BSAShaper;
                classifiers {
                    ieee-802.1 L2BSA vlan-tag outer;
                }
                rewrite-rules {
                    ieee-802.1 L2BSA vlan-tag outer;
                }
            }
        }
        interface-set "$junos-aggregation-interface-set-name" {
            output-traffic-control-profile L2iflsetShaper;
        }
    }
}
```

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, the Layer 2 service support with the pseudowire service logical interfaces is extended to pseudowire service interfaces anchored over redundant logical tunnel interfaces as well.
18.4R1	Starting in Junos OS Release 18.4R1, you can configure the router to detect a logical intermediate node in an access network.
17.1R1	Starting in Junos OS Release 17.1R1, the pseudowire logical tunnel interfaces support Ethernet VPLS, Ethernet bridge, VLAN VPLS, and VLAN bridge encapsulation next hops to exit Layer 2 traffic.

RELATED DOCUMENTATION

Subscriber Management Overview

MPLS Pseudowire Subscriber Logical Interfaces | **331**

Juniper Networks VSAs Supported by the AAA Service Framework

High Availability for Subscriber Access Networks

IN THIS SECTION

- Unified ISSU for High Availability in Subscriber Access Networks | **58**
- Verifying and Monitoring Subscriber Management Unified ISSU State | **59**
- Graceful Routing Engine Switchover for Subscriber Access Networks | **60**
- Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover | **62**

This topic is a high-level overview of high availability for DHCP, L2TP, and PPP access networks.

Unified ISSU for High Availability in Subscriber Access Networks

A unified in-service software upgrade (unified ISSU) enables you to upgrade between two different Junos OS Releases with no disruption on the control plane and with minimal disruption of traffic. The routers preserves the active subscriber sessions and session services across the upgrade, so that they continue after the upgrade has completed.

The unified ISSU feature supports the PPPoE, DHCP, and L2TP access models for subscriber management. Unified ISSU support for the DHCP and L2TP access models was added in Junos OS Release 14.1.

- For static and dynamic PPPoE access, unified ISSU supports the following:
  - 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

> **NOTE:** 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.

- For DHCP access, unified ISSU supports the following:
  - DHCPv4 local server, DHCPv4 relay, DHCPv6 local server, DHCPv6 relay, and DHCP relay proxy
  - Preservation of accounting, filter, and class-of-service (CoS) statistics for DHCP subscribers on MPC/MIC interfaces on MX Series routers
- For L2TP access, unified ISSU supports both the LAC and the LNS. 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.

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.

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 queiscing 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
```

Graceful Routing Engine Switchover for Subscriber Access Networks

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.

To enable GRES support on MX Series routers, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

DHCP

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 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*.

L2TP

GRES is supported on MX Series routers acting as either the L2TP LAC or LNS. In the event that L2TP (jl2tpd, the L2TP universal edge process) restarts or that the router fails over from the active routing engine (RE) to the standby RE, L2TP GRES 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
```

Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover

During a graceful Routing Engine switchover (GRES), access routes and access-internal routes for DHCP and PPP subscriber management can become stale. After the GRES, the router removes any such stale routes from the forwarding table. Some traffic is lost if the stale routes are removed before the routes are reinstalled.

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.

In subscriber networks with nonstop active routing (NSR) 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.

You can reduce the risk of this traffic loss by configuring the router to delay the removal of stale routes after a GRES. The delay period is a nonconfigurable 180 seconds (3 minutes). The router retains the stale routes for the duration of the period, which is long enough 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. The risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

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
   ```

Release History Table

Release	Description
14.1	Unified ISSU support for the DHCP and L2TP access models was added in Junos OS Release 14.1.

RELATED DOCUMENTATION

Getting Started with Unified In-Service Software Upgrade

Routes for DHCP and PPP Subscriber Access Networks | 63

Routes for DHCP and PPP Subscriber Access Networks

IN THIS SECTION

- Access and Access-Internal Routes for Subscriber Management | 63
- Configuring Dynamic Access Routes for Subscriber Management | 64
- Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers | 66
- Suppressing DHCP Access, Access-Internal, and Destination Routes | 68
- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default | 68
- Verifying the Configuration of Access and Access-Internal Routes for DHCP and PPP Subscribers | 69

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-nexthop, automatically resolves to the subscriber's IP address. If you configure the **access-internal** statement in the dynamic profile, it is ignored.

> **NOTE:** Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

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.

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

You can dynamically configure access-internal routes. In releases earlier than Junos OS 15.1, this configuration is optional; if you include it, 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. For PPP subscriber interfaces, you do not need to specify the MAC address for access-internal routes.

To dynamically configure access-internal routes for DHCP or PPP subscribers:

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# set 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. (DHCP subscriber interfaces only) 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
```

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 the following routes:

- DHCPv4 sessions—access-internal and destination routes.
- DHCPv6 sessions—access-internal and access routes.

An access route represents a network behind an attached video services router, and is set to a preference of 13.

An access internal route is a /32 route that represents a directly attached end user, and is set to a preference of 12.

These routes are used by the DHCP application on a video services router to represent either the end users or the networks behind the attached video services router.

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.

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

You can use the route suppression option to override the default route installation behavior. You can configure route suppression and prevent DHCP from installing specific types of routes for:

- DHCP local server and DHCP relay agent
- DHCPv4 and DHCPv6 sessions
- 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.

Example:

- 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
```

Note the following while configuring route suppression option:

- 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.
- The **no-arp** statement supported in legacy DHCP is replaced by the **route-suppression** statement.

Verifying the Configuration of Access and Access-Internal Routes for DHCP and PPP Subscribers

Purpose

View configuration information for access routes and access-internal routes on DHCP and 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
```

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

RELATED DOCUMENTATION

RADIUS IETF Attributes Supported by the AAA Service Framework

DHCP Overview

DHCPv6 Local Server

DHCPv6 Relay Agent DHCPv6 Relay Agent Overview

Subscribers with Identical Framed Routes

Subscribers in the same routing instance are typically expected to have different framed routes. However, there is an active/backup use case where you might want to configure the same framed route for two subscribers. In this scenario, a subscriber expects to receive ingress traffic from an active CPE device, but wants to switch as soon as possible to a backup CPE device.

You can meet this requirement by having two subscribers connected to the same BNG with the same access route for an identical Framed-Route address. However, you must configure a distance value in the

Framed-Route that is different between the two subscribers. For example, you might configure RADIUS as follows:

```
user1@test.com Cleartext-Password := "$abc123"
  Framed-Route = "10.0.0.1/32  distance 12",
  ERX-Virtual-Router-Name  = test,
user2@test.com Cleartext-Password := "$abc123"
  Framed-Route = "10.0.0.1/32  distance 240",
  ERX-Virtual-Router-Name  = test,
```

Subscribers user1 and user2 have the same password, routing instance, and Framed-Route address. The distance is just an administrative distance or preference for discrimination between the routes. The distance is 12 for user1 and 240 for user2. The router can add only one route to the forwarding table. It selects the route with the lowest distance value, which is 12. Consequently, traffic towards the subscriber travels to the logical interface associated with user1.

The router installs the backup route for user2 in the routing table. If the link to user1 goes down, then the router installs the backup route for user2 in the forwarding table so downstream traffic can continue to the subscriber.

What happens if you do not configure different distance values for the two subscribers? Consider the following RADIUS configuration:

```
user1@test.com Cleartext-Password := "$abc123"
  Framed-Route = "10.0.0.1/32",
  ERX-Virtual-Router-Name  = test,
user2@test.com Cleartext-Password := "$abc123"
  Framed-Route = "10.0.0.1/32",
  ERX-Virtual-Router-Name  = test,
```

If both of these subscribers try to log in, only the one that logs in first achieves the Active state, so only that route is installed in the forwarding table. The other subscriber flaps between the Init and Terminated states and never succeeds at logging in, as long as the first subscriber is Active.

2

CHAPTER

DHCP Subscriber Access Networks

DHCP Subscriber Access Networks Overview | **73**

DHCP Snooping for Network Security | **81**

DHCPv4 Duplicate Client Management | **104**

DHCPv6 Duplicate Client Management | **110**

DHCP Subscriber Access Networks Overview

IN THIS SECTION

- DHCP and Subscriber Management Overview | 73
- Subscriber Access Operation Flow Using DHCP Relay | 74
- Defining Various Levels of Services for DHCP Subscribers | 75
- Example: Configuring a Tiered Service Profile for Subscriber Access | 76

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.

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 8 on page 75 shows the flow of operations that occurs when the router is using DHCP relay to enable access for a subscriber.

Figure 8: Subscriber Access Operation Flow

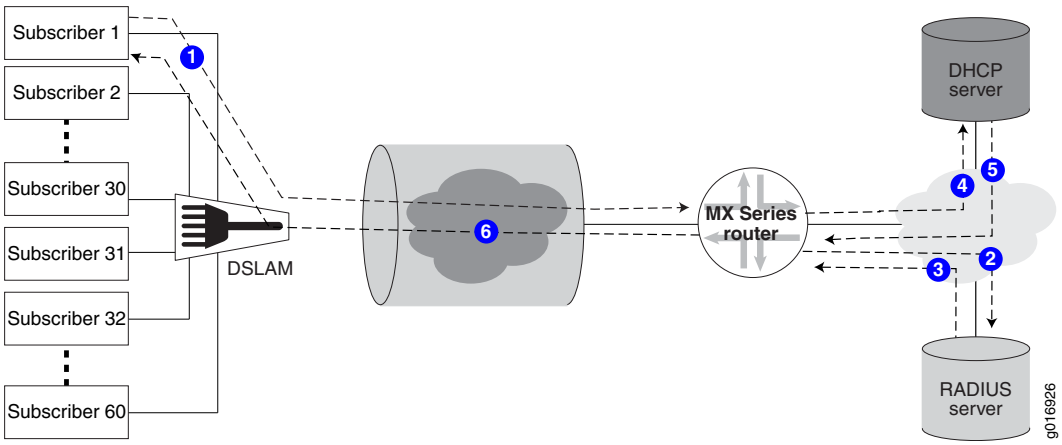

The following general sequence occurs during access configuration for a DHCP client:

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.

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 *RADIUS Servers and 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.

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

- *DHCP Overview*
- *DHCPv6 Local Server*
- *DHCPv6 Relay Agent*
- *Address-Assignment Pools for Subscriber Management*
- *Introduction to Subscriber Management*
- *Dynamic Profiles for Subscriber Management*
- *CoS for Subscriber Access Overview*

DHCP Snooping for Network Security

IN THIS SECTION

- DHCP Snooping Support | 81
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server | 83
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent | 84
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent | 91
- Disabling DHCP Snooping Filters | 93
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent | 95
- Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent | 97
- Preventing DHCP Spoofing | 103

DHCP Snooping Support

What is DHCP Snooping

DHCP snooping provides additional security by identifying the incoming DHCP packets and rejecting DHCP traffic determined to be unacceptable from untrusted devices in the network.

DHCP allocates IP addresses dynamically, leasing addresses to devices so that the addresses can be reused when they are no longer needed by the devices to which they were assigned. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping looks into incoming DHCP packets and examines DHCP messages. It extracts their IP addresses and lease information allocated to clients and builds up a database. Using this database, it can determine if the packets arriving are from the valid clients—that is—the IP addresses of the clients was assigned by the DHCP server. In this way, DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port).

Benefits of DHCP Snooping

- DHCP snooping provides an extra layer of security via dynamic IP source filtering.
- DHCP snooping can prevent rogue DHCP activity in the network by filtering out DHCP packets that are arriving on the wrong ports, or with incorrect contents.

Configuring DHCP Snooping

In the default DHCP snooping configuration, all traffic is snooped. 3146531465

On Junos OS device, DHCP snooping is enabled in a routing instance when you configure the following options in that routing instance:

- **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level
- **dhcp-local-server** statement at the **[edit system services]** hierarchy level
- 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.

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.

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 83 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;
        }
    }
}
```

SEE ALSO

DHCP Snooping Support

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*, 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;
                    }
                }
            }
        }
    }
}
```

SEE ALSO

DHCP Snooping Support

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

Example: Configuring DHCP Snooping Support for DHCP Relay Agent | **95**

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, you enable or disable the DHCP relay snooping feature.

Table 5 on page 91 shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the **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 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
forward-snooped-clients not configured	snooped packets result in subscriber (DHCP client) creation	dropped
all-interfaces	forwarded	forwarded
configured-interfaces	forwarded	dropped
non-configured-interfaces	snooped packets result in subscriber (DHCP client) creation	forwarded

Table 6 on page 91 shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the **no-allow-snooped-clients** statement.

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
forward-snooped-clients not configured	dropped	dropped
all-interfaces	dropped	forwarded

Table 6: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled *(continued)*

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
configured-interfaces	dropped	dropped
non-configured-interfaces	dropped	forwarded

Table 7 on page 92 shows the action the router (or switch) takes for the snooped BOOTREPLY packets.

Table 7: Actions for Snooped BOOTREPLY Packets

forward-snooped-clients Configuration	Action
forward-snooped-clients not configured	snooped **BOOTREPLY** packets dropped if client is not found
forward-snooped-clients all configurations	snooped **BOOTREPLY** packets forwarded if client is not found

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.

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;
    }
}
```

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
```

SEE ALSO

DHCP Snooping Support

no-snoop | 677

Example: Configuring DHCP Snooping Support for DHCP Relay Agent

IN THIS SECTION

- Requirements | 95
- Overview | 95
- Configuration | 95

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

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
```

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

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

3. 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
```

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

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

5. Enable DHCP snooping support for the group.

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

6. 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
```

7. Enable DHCP snooped packet forwarding for DHCP relay agent.

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

8. 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.

Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent

IN THIS SECTION

- Requirements | **98**
- Overview | **98**
- Configuration | **98**
- Verification | **101**

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

This example uses the following hardware and software components:

- MX Series 5G Universal Routing Platforms
- 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

IN THIS SECTION

- Enabling DHCPv6 Snooping Support for a Named Group of Interfaces | **99**
- Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group | **100**

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:

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

IN THIS SECTION

- Verifying the Address Bindings for DHCPv6 Relay Agent Clients | **101**

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

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:                         LL0x1-00:00:5e:00:53:02
     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.

SEE ALSO

DHCPv6 Relay Agent Overview

DHCP Snooping Support

Grouping Interfaces with Common DHCP Configurations

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent | 84

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 84.

SEE ALSO

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 | 84

DHCPv4 Duplicate Client Management

IN THIS SECTION

- DHCPv4 Duplicate Client In Subnet Overview | 105
- Guidelines for Configuring Support for DHCPv4 Duplicate Clients | 105
- Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information | 106
- Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces | 108

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.

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.

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
```

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
```

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

DHCPv6 Duplicate Client Management | 110

DHCPv6 Duplicate Client Management

IN THIS SECTION

- DHCPv6 Duplicate Client DUIDs | **110**
- Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs | **110**

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.

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
```

Release History Table

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

RELATED DOCUMENTATION

DHCPv4 Duplicate Client Management | 104

3

CHAPTER

PPP Subscriber Access Networks

PPP Subscriber Access Networks Overview | **114**

PPP Network Control Protocol Negotiation | **136**

Tracing PPP Service Events for Troubleshooting | **145**

PPP Subscriber Access Networks Overview

IN THIS SECTION

- Dynamic Profiles for PPP Subscriber Interfaces Overview | 114
- Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests | 115
- RADIUS-Sourced Connection Status Updates to CPE Devices | 118
- Configuring Dynamic Profiles for PPP | 123
- Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges | 124
- How to Configure RADIUS-Sourced Connection Status Updates to CPE Devices | 125
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | 126
- Migrating Static PPP Subscriber Configurations to Dynamic Profiles Overview | 127
- Configuring Local Authentication in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | 128
- Configuring Tag2 Attributes in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | 130
- Configuring Dynamic Authentication for PPP Subscribers | 131
- Modifying the CHAP Challenge Length | 133
- Example: Minimum PPPoE Dynamic Profile | 134
- Verifying and Managing PPP Configuration for Subscriber Management | 135

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.

Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests

IN THIS SECTION

- Benefits of PPP Fast Keepalives | **115**
- How PPP Fast Keepalive Processing Works | **116**
- Statistics Display for PPP Fast Keepalive | **116**
- Effect of Changing the Forwarding Class Configuration | **117**
- Ignoring a Magic Number Mismatch | **117**

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

- 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.

Ignoring a Magic Number Mismatch

When the Packet Forwarding Engine validates the peer magic number in the received LCP Echo-Request 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.

SEE ALSO

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 | **124**

RADIUS-Sourced Connection Status Updates to CPE Devices

Starting in Junos OS Release 20.2R1, you can use RADIUS-sourced messages to convey information that the BNG transparently forwards to a CPE device, such as a home gateway. For example, this information might be upstream bandwidth or some other connection rate parameter that the CPE device needs. This capability is useful when you want to dynamically enforce traffic management as close to subscribers as possible.

Ordinarily, you might use the RADIUS standard attribute Reply-Message (18) to convey this information to the CPE device during PPP authentication. However, if you are already using that attribute for something else, you can also use the Juniper Networks Connection-Status-Message VSA (26-4874–218). This VSA is a logical extension to the Reply-Message attribute (18) and has the same format and semantics.

PPP uses a Juniper Networks vendor-specific extension to LCP to send the contents of the Connection-Status-Message VSA to the peer home gateway. PPP includes this information in the Connection-Status-Message option of an LCP Connection-Update-Request message.

RADIUS can send the Connection-Status-Message VSA to authd in the following ways:

- In the RADIUS Access-Accept message during negotiation and authorization of a PPP session
- In a RADIUS CoA request at any time for an active PPP session

You might use both of these methods for any given session for business or residential subscribers. The Access-Accept message provides the initial connection parameters. The CoA capability enables you to update connection rate parameters as needed throughout the life of a session. The information carried in the Connection-Status-Message VSA is typically traffic rates that are applied by local configuration such as a dynamic service profile or the corresponding ANCP Port Up message.

> **NOTE:** If you do not include the **lcp-connection-update** PPP option in the dynamic client profile, PPP processes the notification from authd, but takes no action. If LCP on the router is not in the Opened state, then PPP takes no action on the VSA.

The following steps describe what happens when RADIUS sends the VSA in an Access-Accept message:

1. The authd process receives the Connection-Status-Message VSA in an Access-Accept message from the RADIUS server.
2. The authd process sends the Connection-Status-Message VSA to PPP (jpppd).
3. PPP NCP negotiation takes place between the remote gateway PPP client and PPP on the router.

4. Successful negotiation results in a family activation request. The PPP session enters the Session Up state when the family is activated.
5. If the dynamic client profile includes the **lcp-connection-update** PPP option and LCP on the router is in the Opened state, PPP sends an LCP Connection-Update-Request message to the gateway. This message includes the VSA information in the Connection-Status-Message option.
   - If the gateway supports the LCP Connection-Update-Request, it returns an LCP Connection-Update-Ack message to the router. The home gateway LCP must be in the Opened state when it receives the request, otherwise it discards the request.
   - If the gateway does not support the LCP Connection-Update-Request, it returns an LCP Code-Reject message to the router.

> **NOTE:** If the gateway does not respond, the router retries the update request. It uses the PPP default values of up to a maximum of 10 retries with an interval of 3 seconds between the attempts.

> **NOTE:** If you do not include the **lcp-connection-update** PPP option in the dynamic client profile, PPP processes the notification from authd, but takes no action. If the option is present but LCP on the router is not in the Opened state, PPP takes no action regarding the VSA.

The following steps describe what happens when RADIUS sends the VSA in a CoA request. This assumes that NCP negotiation was already successful and the session is active.

1. The authd process receives the Connection-Status-Message VSA in a CoA request from the RADIUS server.
2. The authd process sends the Connection-Status-Message VSA to PPP (jpppd).
3. If the dynamic client profile includes the **lcp-connection-update** PPP option and LCP on the router is in the Opened state, PPP sends an LCP Connection-Update-Request message to the gateway. This message includes the VSA information in the Connection-Status-Message option.
   - If the gateway supports the LCP Connection-Update-Request, it returns an LCP Connection-Update-Ack message to the router. The home gateway LCP must be in the Opened state when it receives the request, otherwise it discards the request.
   - If the gateway does not support the LCP Connection-Update-Request, it returns an LCP Code-Reject message to the router.

> **NOTE:** If the gateway does not respond, the router retries the update request. It uses the PPP default values of up to a maximum of 10 retries with an interval of 3 seconds between the attempts.

If the home gateway fails to receive a Connection-Update-Request message, the router retries sending the message. The router also retries the request when it does not receive either a Connection-Update-Ack or an LCP Code-Reject back from the gateway, or when something is wrong with the Ack message. The default retry interval is 3 seconds. The router will retry the message up to the default 10 times before it quits. If the router exhausts all the retry attempts without receiving an appropriate Connection-Update-Ack message, it logs the message as if it had received a PPP Code-Reject message.

> **NOTE:** RADIUS can include multiple instances of the Connection-Status-Message VSA in either the Access-Accept message or a CoA request. If this occurs, authd uses only the first instance and ignores any others.

The Access-Accept or CoA requests might contain other attributes besides the Connection-Status-Message VSA, but there is no interdependency between the VSA and any other attributes. This is true even when the message includes the Activate-Service (26–65) or Deactivate-Service (26–66) VSAs. The lack of dependency means that even if authd does not successfully apply the other attributes, it still sends the connection info to PPP, which in turn sends the VSA contents to the home gateway.

Similarly, authd applies any other attributes and returns a CoA response regardless of whether PPP successfully communicates the content of the Connection-Status-Message VSA to the remote gateway. This is true even when the CoA contains only the Connection-Status-Message VSA. This capability is necessary because not all gateways will accept the LCP extension used in this feature.

Message and Option Formats

Figure 9 on page 121 shows the format for Connection-Update-Request and Connection-Update-Ack messages. The formats are the same, but Table 8 on page 121shows that some field values are different for the two messages.

Figure 9: Connection-Update-Request and Connection-Update-Ack Message Format

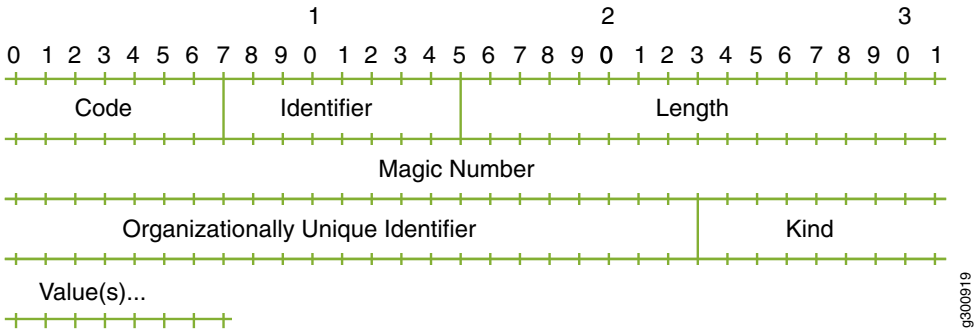

Table 8: Field Values for Connection-Update-Request and Connection-Update-Ack messages

Field	Connection-Update-Request	Connection-Update-Ack
Code	0 for vendor-specific	0 for vendor-specific
Identifier	Identifier for vendor-specific packet	Same identifier as in the Connection-Update-Request message. If this value does not match, the router logs the error and discards the packet. This enables the request message to be retried, just as if the gateway had not received it.
Length	Number of bytes in the packet: 12 plus the length of the Connection-Status-Message option	Number of bytes in the Connection-Update-Ack packet: 12
Magic Number	Negotiated value for the local PPP magic number	Negotiated value for the local PPP magic number
Organizationally Unique Identifier (OUI)	00-21-59 for Juniper Networks	00-21-59 for Juniper Networks
Kind	1 for Session-Update	2 for Session-Ack. For any other value, the router logs the error and the discards the packet. This enables the request message to be retried, just as if the gateway had not received it.
Values	Connection-Status-Message option in TLV format	No values are supported

You can configure how the PPP magic numbers are used.

- If you configure **ignore-magic-number-mismatch** PPP option, you are preventing the magic numbers from being validated. PPP ignores a mismatch between the magic numbers in the request and Ack messages. If there are no other validation errors, PPP accepts the Connection-Update-Ack message.
- If you do not configure **ignore-magic-number-mismatch** PPP option, the magic numbers go through validation. If the magic number in the Ack message does not match the gateway's magic number established during LCP negotiation, the router logs the error and discards the Connection-Update-Ack message as an invalid response. This enables the request message to be retried, just as if the gateway had not received it.

See "Preventing the Validation of PPP Magic Numbers During PPP Keepalive Exchanges" on page 124 for more information about magic numbers.

Figure 10 on page 122 shows the format for the Connection-Status-Message options. Table 9 on page 122 lists the field values.

Figure 10: Connection-Status-Message Option Format

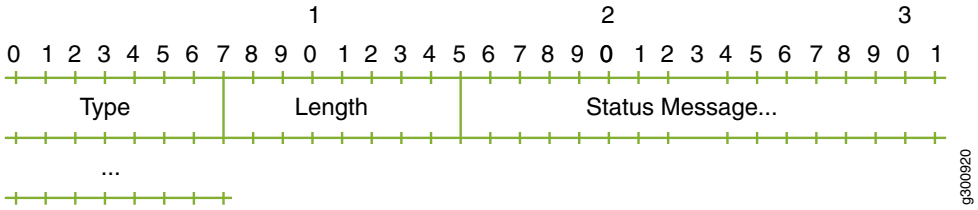

Table 9: Field Values for the Connection-Status-Message Option

Field	Value
Type	1
Length	Number of bytes in the option; 2 plus the length of the message. The message length can be from 1 through 247 bytes.
Status Message	Contents of the Connection-Status-Message VSA

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 126 in the *Junos Subscriber Access Configuration Guide*.

For information about using dynamic profiles to authenticate PPP subscribers, see "Configuring Dynamic Authentication for PPP Subscribers" on page 131.

NOTE: Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces for this release.

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.

How to Configure RADIUS-Sourced Connection Status Updates to CPE Devices

You can use RADIUS-sourced messages to convey information that the BNG transparently forwards to a CPE device, such as a home gateway. For example, this information might be upstream bandwidth or some other connection rate parameter that the CPE device needs.

When you enable this feature, PPP can act on a Connection-Status-Message VSA (26–218) received by authd in either a RADIUS Access-Accept message or a CoA message. PPP then conveys the contents of the VSA in an LCP Connection-Update-Request message to the remote peer. This action requires the following to be true:

- At least the first address family has been successfully negotiated and the session is active.
- The router LCP is in the Opened state.

Otherwise PPP takes no action on the VSA. If you do not enable the **lcp-connection-update** option, PPP processes the notification from authd, but takes no action.

You configure this capability in the dynamic client profile associated with subscribers using the CPE device. In practice, you are adding this to numerous other features in the client profile. This example shows only the specific configuration for this feature. This feature also requires you to configure VSA 26-218 on your RADIUS server; that is outside the scope of this documentation.

To configure the connection status updates in a dynamic profile for PPP subscriber interfaces:

1. Edit the PPP options in the client profile.

```
[edit dynamic-profiles ppp-client-profile interfaces pp0 unit “$junos-interface-unit”]
user@host# edit dynamic-profiles ppp-client-profile interfaces pp0 unit “$junos-interface-unit” ppp-options
```

2. Enable the connection status updates.

```
[edit dynamic-profiles ppp-client-profile interfaces pp0 unit “$junos-interface-unit” ppp-options]
user@host# set lcp-connection-update
```

You can use the **show ppp interface extensive** command for the PPP logical interface to determine whether LCP connection updates are successful. You can monitor the relevant statistics with the **show system subscriber-management statistics ppp** command.

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
```

Migrating Static PPP Subscriber Configurations to Dynamic Profiles Overview

This topic discusses several considerations for migrating certain static, terminated IPv4 PPP subscriber configurations to dynamic configurations using dynamic profiles. Service providers managing static subscribers on routers with legacy Junos OS releases (earlier than Junos OS Release 15.1R4) have requirements for migrating their static subscribers to being managed with dynamic profiles on routers running enhanced subscriber management (Junos OS Release 15.1R4 and later releases). Starting in Junos OS Release 18.2R1, several enhancements have been added to facilitate the transition of these static service provider configurations to dynamic profiles.

Local Authentication

Some providers with static configurations might use CPE devices that do not support any authentication protocols, not even CHAP or PAP. The providers might use PPPoE service name tables as a rudimentary method to authenticate and authorize the subscribers on static PPPoE logical interfaces. If the subscriber ACI or ARI do not match a table entry, then the PPP PADI and PADR packets are typically dropped. Legacy Junos OS releases do not support subscribers configured with *no-authentication* for the authentication method.

For subscribers where the CPE does not support authentication protocols such as PAP and CHAP, you can configure usernames and passwords locally. The router uses these values when it contacts the RADIUS server for authentication.

- To configure the username for local authentication, include the **username-include** statement in the PPP options for the dynamic logical interface. You can define the name based on one or more of the following attributes: MAC address, Agent Circuit ID, Agent Remote ID, and domain name. By default, a period (.) is the delimiter between elements of the name, but you can define other characters instead.
- To configure the password for local authentication, include the **password** statement in the PPP options for the dynamic logical interface.

You can use the same dynamic profile to support both CPEs that do not support an authentication protocol and CPEs that do.

CPE-Sourced Address Assignment

For some static configurations, the subscriber address is not assigned by using RADIUS or a local address pool on the router. Instead, the CPE has a static address configured for the subscriber; during IPCP negotiation, the CPE requests the router to assign that address to the subscriber.

Starting in Junos OS Release 18.2R1, you can assign a wildcard address of 255.255.255.255 to the Framed-Route-Address attribute [8] in the configuration for your RADIUS server. When RADIUS returns

the attribute with that value, jpppd automatically accepts the subscriber IP address assignment provided by the client in an IPCP configure-request message rather than assigning another address.

Tag2 Route Attribute

In some configurations, static PPP subscriber interfaces are configured in different VRFs. Each VRF configuration has static routes that point to static PPP subscriber interfaces as the next-hop address. These routes might have the tag2 attribute configured; it is required by MP-BGP to apply the appropriate local preference and community when it advertises the routes.

Starting in Junos OS Release 18.2R1, you can configure your RADIUS server to include the tag2 attribute in the Framed-Route attribute [22] when it authenticates a subscriber.

You must also configure the dynamic profile to derive the tag2 value from the Framed-Route attribute. To do so, specify the $junos-framed-route-tag2 predefined variable to be used when the access route is dynamically instantiated. Alternatively, you can configure the dynamic profile to provide a specific tag2 value for a specific access route prefix.

See *Junos OS Predefined Variables* for more information about predefined variables.

Benefits

- Local authentication enables authentication with locally stored passwords and usernames for subscribers when the CPE does not support authentication protocols such as PAP and CHAP.
- CPE-sourced address assignment enables the router to accept statically configured subscriber IP addresses requested by the CPE rather than assigning the address from a local or externally sourced address pool.
- The tag2 attribute enables more detailed specification of routes.

Configuring Local Authentication in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers

Some providers with static configurations might use CPE devices that do not support any authentication protocols, not even CHAP or PAP. The providers might use PPPoE service name tables as a rudimentary method to authenticate and authorize the subscribers on static PPPoE logical interfaces. If the subscriber ACI or ARI does not match a table entry, then the PPP PADI and PADR packets are typically dropped.

Starting in Junos OS Release 18.2R1, you can configure usernames and passwords locally for clients that do not support authentication protocols such as PAP and CHAP. The router uses these values when it contacts the RADIUS server for authentication. This aids in the migration of the static subscribers to using dynamic profiles on a router running enhanced subscriber management.

To configure local authentication:

1. Configure the username using one or more of the available options.

   a. (Optional) Specify that the agent circuit identifier (ACI) is included in the username. The ACI is derived from PPPoE tags.

      ```
      [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
        ppp-options local-authentication]
      user@host# set username-include circuit-id
      ```

   b. (Optional) Specify that the agent remote ID (ARI) is included in the username. The ARI is derived from PPPoE tags.

      ```
      [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
        ppp-options local-authentication]
      user@host# set username-include remote-id
      ```

   c. (Optional) Specify that the MAC address from the client PDU is included in the username. The MAC address is derived from the PPPoE packet.

      ```
      [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
        ppp-options local-authentication]
      user@host# set username-include mac-address
      ```

   d. (Optional) Specify the client domain name to end the username. The router adds the @ character as the delimiter for this option.

      ```
      [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
        ppp-options local-authentication]
      user@host# set username-include domain-name name
      ```

   e. (Optional) Specify a delimiter to separate the components that make up the username. The default delimiter is a period (.).The router always uses the @ character as the delimiter before the domain name.

      ```
      [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
        ppp-options local-authentication]
      user@host# set username-include delimiter character
      ```

2. Configure the password for the subscriber.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
  ppp-options local-authentication]
user@host# set password password
```

The username takes the following format when you include all the options and use the default delimiter:

```
mac-address.circuit-id.remote-id@domain-name
```

For example, consider the following sample configuration, where the ACI is aci1002, the ARI is ari349, and the MAC address is 00:00:5e:00:53:ff:

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options
  local-authentication]
user@host# set username-include circuit-id
user@host# set username-include remote-id
user@host# set username-include mac-address
user@host# set username-include domain-name example.com
user@host# set username-include delimiter -
user@host# set password $ABC123$ABC123
```

This configuration results in a local password of $ABC123$ABC123 for the following unique local username:

0000.5e00.53ff-aci1002-ari349@example.com

Configuring Tag2 Attributes in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers

In some configurations, PPP subscribers use static routes with a tag2 attribute. For example, MP-BGP uses tag2 to enable it to apply the appropriate local-preference and community when it advertises routes. When you migrate these subscribers to using dynamic profiles on a router running enhanced subscriber management, you can configure the tag2 attribute by configuring a specific value for a route or by deriving the value from a RADIUS server. This support is first available in Junos OS Release 18.2R1.

- To configure a specific tag2 value for a route:
  - Specify the value.

    ```
    [edit dynamic-profiles profile-name routing-options access route prefix]
    user@host# set tag2 route-tag2
    ```

- To derive the tag2 value from a RADIUS server:

  1. Configure your RADIUS server to include the tag2 attribute in the Framed-Route attribute [22] when it authenticates a subscriber. Consult your RADIUS server documentation for configuration information. The configuration might look something like the following example:

     ```
     user@sub.example.com User-Password := "$ABC123"
        Service-Type = Framed-User,
        Framed-Protocol = PPP,
        Framed-Route += "198.51.100.0/24 203.0.113.27 tag 5 distance 10 tag2 3"
     ```

  2. Configure the dynamic profile to use the $junos-framed-route-tag2 predefined variable to dynamically derive the tag2 value from the Framed-Route attribute.

     ```
     [edit dynamic-profiles profile-name routing-options access route "$junos-framed-route-ip-address-prefix]
     user@host# set tag2 $junos-framed-route-tag2
     ```

     The $junos-framed-route-ip-address-prefix predefined variable derives the IPv4 address prefix for the access route from the Framed-Route attribute as well.

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 133.

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 140.

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
```

   See "Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses" on page 144 for more information.

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 131.
  - 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
```

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";
            }
        }
    }
}
```

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
```

Release History Table

Release	Description
20.2R1	Starting in Junos OS Release 20.2R1, you can use RADIUS-sourced messages to convey information that the BNG transparently forwards to a CPE device, such as a home gateway.
18.2R1	Starting in Junos OS Release 18.2R1, several enhancements have been added to facilitate the transition of these static service provider configurations to dynamic profiles.
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 | 124

Dynamic Profiles Overview

Configuring a Basic Dynamic Profile

Junos OS Predefined Variables

For L2TP group profiles: Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

PPP Network Control Protocol Negotiation

IN THIS SECTION

- PPP Network Control Protocol Negotiation Mode Overview | 136
- Controlling the Negotiation Order of PPP Authentication Protocols | 140
- Configuring the PPP Network Control Protocol Negotiation Mode | 142
- Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses | 144

PPP Network Control Protocol Negotiation Mode Overview

IN THIS SECTION

- PPP NCP Negotiation Modes | 137
- PPP NCP Negotiation Mode Supported Configurations | 138
- PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers | 138
- PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers | 138
- PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations | 139

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.

Junos OS supports the following NCPs as presented in the associated IETF standards:

- Internet Protocol Control Protocol (IPCP) in RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
- IPv6 Control Protocol (IPv6CP) in RFC 5072, *IP Version 6 over PPP*

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 10 on page 137.

Table 10: 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 adress 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 adress 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.

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 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 131.
  - 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*.
  - For information about dynamic profiles for PPP subscribers, see "Dynamic Profiles for PPP Subscriber Interfaces Overview" on page 114.

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]
  ```

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
     ```

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 ensures that the CPE can use 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.
14.1	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.
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

Configuring the PPP Attributes for a Group Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers | **131**

Tracing PPP Service Events for Troubleshooting

IN THIS SECTION

- Configuring the PPP Service Trace Log Filename | **147**
- Configuring the Number and Size of PPP Service Log Files | **147**
- Configuring Access to the PPP Service Log File | **148**
- Configuring a Regular Expression for PPP Service Messages to Be Logged | **148**

- Configuring Subscriber Filtering for PPP Service Trace Operations | **149**
- Configuring the PPP Service Tracing Flags | **150**
- Configuring the Severity Level to Filter Which PPP Service Messages Are Logged | **150**

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 147.

2. (Optional) Configure the number and size of trace logs.

   See "Configuring the Number and Size of PPP Service Log Files" on page 147.

3. (Optional) Configure user access to trace logs.

   See "Configuring Access to the PPP Service Log File" on page 148.

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 148.

5. (Optional) Configure flags to specify which events are logged.

   See "Configuring the PPP Service Tracing Flags" on page 150.

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 150.

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
```

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
```

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
```

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
```

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
```

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
```

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

PPP Subscriber Access Networks Overview | 114

4

CHAPTER

L2TP Subscriber Access Networks

L2TP for Subscriber Access Overview | **153**

L2TP Tunnel Switching For Multiple-Domain Networks | **166**

L2TP LAC Subscriber Configuration | **183**

L2TP LAC Tunneling for Subscribers | **191**

L2TP Subscriber Access Lines and Connection Speeds | **225**

L2TP LNS Inline Service Interfaces | **260**

IP Packet Reassembly on Inline Service Interfaces | **317**

Peer Resynchronization After an L2TP Failover | **320**

Tracing L2TP Events for Troubleshooting | **323**

L2TP for Subscriber Access Overview

IN THIS SECTION

- L2TP for Subscriber Access Overview | 153
- L2TP Terminology | 156
- L2TP Implementation | 157
- Retransmission of L2TP Control Messages | 159
- Configuring Retransmission Attributes for L2TP Control Messages | 160
- Enabling Tunnel and Global Counters for SNMP Statistics Collection | 163
- Verifying and Managing L2TP for Subscriber Access | 164

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 11 on page 153 shows a simple L2TP topology.

Figure 11: Typical L2TP Topology

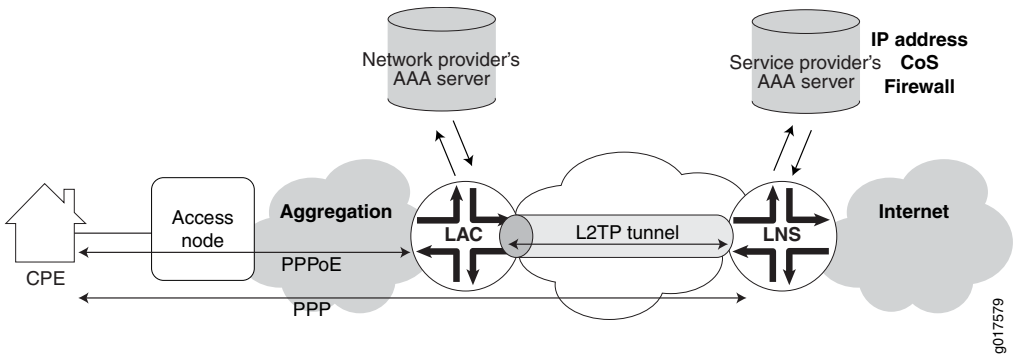

L2TP separates the termination of access technologies, such as cable or xDSL, from the termination of PPP and subsequent access to a network. This separation enables public ISPs to outsource their access

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 12 on page 154 shows the protocol layer stacking for an L2TP pass-through connection.

Figure 12: Protocol Stacking for L2TP Subscribers in Pass-Through Mode

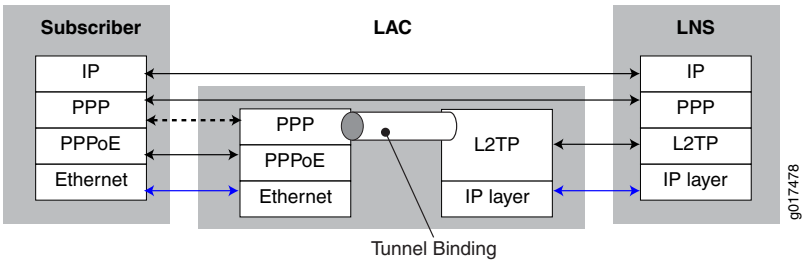

NOTE: On MX Series routers, the LAC and LNS functions are supported only on MPCs; they are not supported on any services PIC or MS-DPC. For details about MPC support for L2TP, see the *MX Series Interface Module Reference*

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.

NOTE: L2TP is not supported over GRE tunnels.

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.

L2TP Terminology

Table 11 on page 156 describes the basic terms for L2TP.

Table 11: 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.

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 SCCRP 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.

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.

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.

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 12 on page 163 have a default value of zero.

Table 12: 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
```

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 | 183

Configuring an L2TP LNS with Inline Service Interfaces | 261

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 | **216**

Domain Mapping Overview

L2TP Tunnel Switching For Multiple-Domain Networks

IN THIS SECTION

- L2TP Tunnel Switching Overview | **166**
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary | **170**
- Configuring L2TP Tunnel Switching | **175**
- Setting the L2TP Receive Window Size | **177**
- Setting the L2TP Tunnel Idle Timeout | **178**
- Setting the L2TP Destruct Timeout | **178**
- Configuring the L2TP Destination Lockout Timeout | **180**
- Removing an L2TP Destination from the Destination Lockout List | **181**
- Configuring L2TP Drain | **181**
- Using the Same L2TP Tunnel for Injection and Duplication of IP Packets | **182**

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 13 on page 167. 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 13 on page 167, 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 13: L2TP Tunnel Switching Network Topology

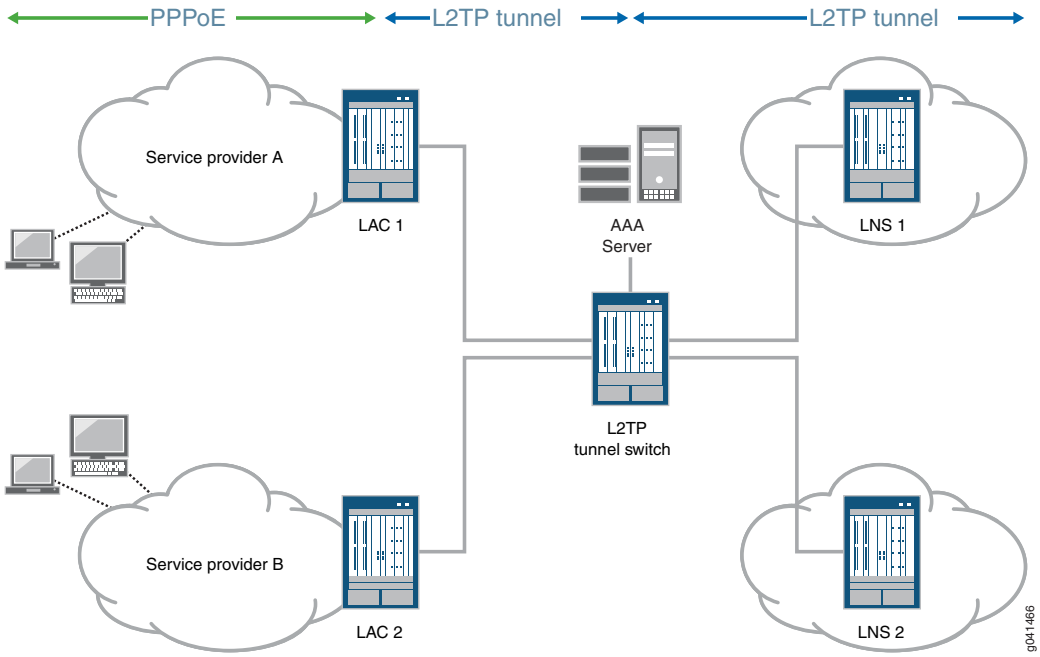

L2TP tunnel switching simplifies network configuration when the administrative domain of a LAC is different from that of the desired LNS. For example:

- 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 14 on page 168 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 14: L2TP Tunnel Switching for Incoming Calls

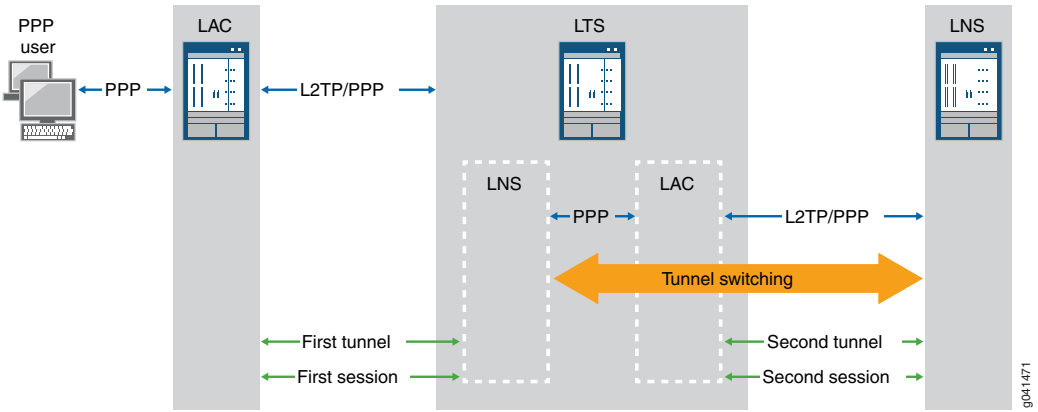

Application of Tunnel Switch Profiles

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 13 on page 169).

Table 13: Cause of CDN Message

CDN Message Is Received On	When
LAC interface	Either of the following occurs: • The second session cannot be established. • The remote LNS terminates the second session.
LNS interface	Either of the following occurs: • 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 14 on page 169).

Table 14: Cause of StopCCN Message

StopCCN Message Is Received On	When
LAC interface	Either of the following occurs: • 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 15 on page 170 lists the actions taken when an administrative **clear** command is issued on the LTS.

Table 15: 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.
clear services l2tp destination all	Clear all destinations and all associated tunnels and sessions.	None.
clear services l2tp session	Clear the session.	Clear the corresponding mapped switched session for this session by sending it a CDN message with the cause set to Administrative.
clear services l2tp session all	Clear all sessions.	None.
clear services l2tp tunnel	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.
clear services l2tp tunnel all	Clear all tunnels.	None.

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 16 on page 171 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 16: 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

Table 16: Default Action for Handling L2TP AVPs at the Switching Boundary *(continued)*

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
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.
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

Table 16: Default Action for Handling L2TP AVPs at the Switching Boundary *(continued)*

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
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.
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.

Table 16: Default Action for Handling L2TP AVPs at the Switching Boundary *(continued)*

AVP Name (Number)	AVP Type	L2TP Message Type	Default Action
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

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.

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
```

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
```

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
```

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.

> **NOTE:** In general, a locked destination cannot be used until the lockout timer expires. However, when L2TP performs the tunnel selection process, in some circumstances it clears the lockout timer for a locked destination. See *Selection When Failover Between Preference Levels Is Configured* and *Selection When Failover Within a Preference Level Is Configured* in "LAC Tunnel Selection Overview" on page 191 for detailed information about the selection process.

> **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.

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
```

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**.

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 form 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

RELATED DOCUMENTATION

- L2TP for Subscriber Access Overview | **153**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**
- *Specifying a Tunnel Switch Profile in a Domain Map*
- LAC Tunnel Selection Overview | **191**

L2TP LAC Subscriber Configuration

IN THIS SECTION

- Configuring an L2TP LAC | **183**
- Configuring How the LAC Responds to Address and Port Changes Requested by the LNS | **186**
- LAC Interoperation with Third-Party LNS Devices | **189**
- Globally Configuring the LAC to Interoperate with Cisco LNS Devices | **189**

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 216.

2. (Optional) Configure the method used for selecting among multiple tunnels.
   - See "Configuring the L2TP LAC Tunnel Selection Parameters" on page 219.
   - See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 221.

- See “Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions” on page 221.
- See “Configuring LAC Tunnel Selection Failover Within a Preference Level” on page 220.

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 252.

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 244.

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 321.

6. (Optional) Specify the format for the tunnel name.

   See “Setting the Format for the Tunnel Name” on page 215.

7. (Optional) Specify when and how many times L2TP retransmits unacknowledged control messages.

   See “Configuring Retransmission Attributes for L2TP Control Messages” on page 160.

8. (Optional) Specify how long a tunnel can remain idle before being torn down.

   See “Setting the L2TP Tunnel Idle Timeout” on page 178.

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 177.

10. (Optional) Specify how long the router retains information about terminated dynamic tunnels, sessions, and destinations.

    See “Setting the L2TP Destruct Timeout” on page 178.

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 186.

12. (Optional) Configure all tunnels on the LAC for interoperation with Cisco LNS devices.

    See “Globally Configuring the LAC to Interoperate with Cisco LNS Devices” on page 189.

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 247.

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 222.

15. (Optional) Prevent the creation of new sessions, destinations, or tunnels for L2TP.

    See “Configuring L2TP Drain” on page 181.

16. (Optional) Enable SNMP statistics counters.

    See “Enabling Tunnel and Global Counters for SNMP Statistics Collection” on page 163.

17. (Optional) Configure trace options for troubleshooting the configuration.

    See “Tracing L2TP Events for Troubleshooting” on page 323.

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
```

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.

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

- L2TP for Subscriber Access Overview | **153**
- Configuring an L2TP LAC | **183**
- Configuring a Tunnel Profile for Subscriber Access | **216**
- *Juniper Networks VSAs Supported by the AAA Service Framework*

L2TP LAC Tunneling for Subscribers

IN THIS SECTION

- LAC Tunnel Selection Overview | 191
- L2TP Session Limits Overview | 207
- Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | 213
- Setting the Format for the Tunnel Name | 215
- Configuring a Tunnel Profile for Subscriber Access | 216
- Configuring the L2TP LAC Tunnel Selection Parameters | 219
- Configuring LAC Tunnel Selection Failover Within a Preference Level | 220
- Configuring Weighted Load Balancing for LAC Tunnel Sessions | 221
- Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions | 221
- Enabling the LAC for IPv6 Services | 222
- Testing L2TP Tunnel Configurations from the LAC | 222

LAC Tunnel Selection Overview

When a user logs in to a domain, the PPP client contacts the LAC 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 194 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 199 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 201 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 202 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 201 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 15 on page 196 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 15: Destination and Tunnel Selection Process with Failover Between Preference Levels

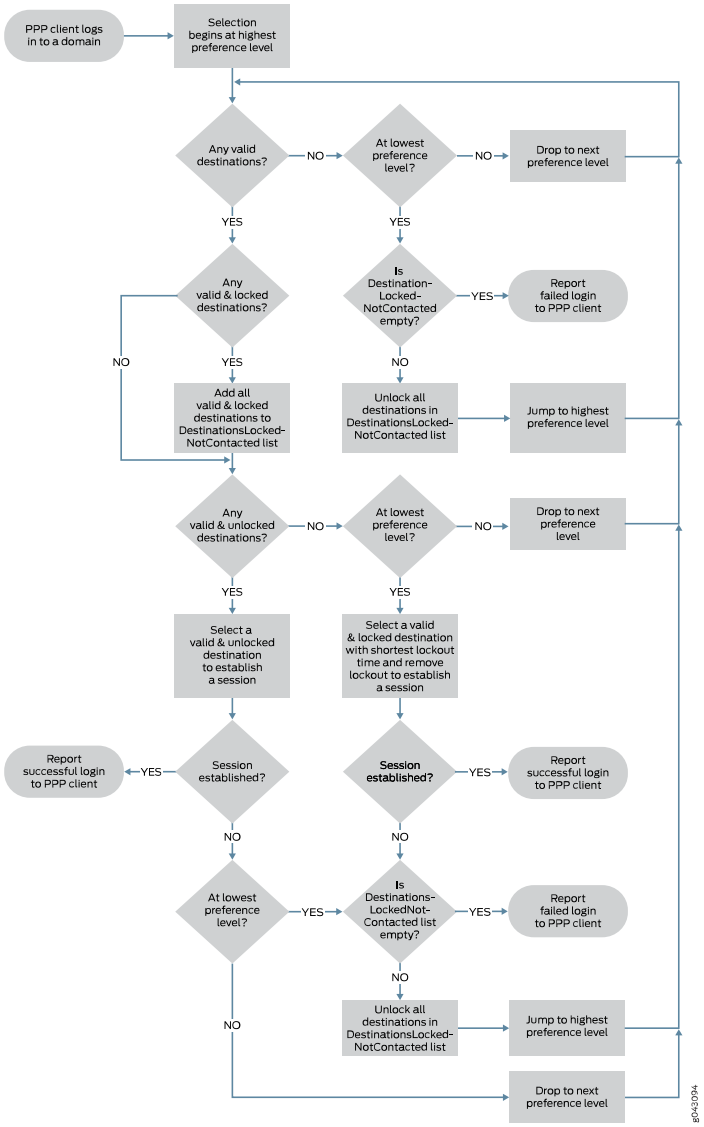

For example, suppose that the tunnel profile includes the following tunnels, each with a valid destination:

- 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.

If the LAC drops down to the lowest preference level and finds no unlocked, valid destinations, the behavior depends on the DestinationsLockedNotContacted list:

- If 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 the DestinationsLockedNotContacted is empty—meaning that all valid destinations have been attempted—then the LAC reports a failed login to the PPP client.

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 it's 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.

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 17 on page 208, 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 17: Scenario 1, Chassis Limit

Level	Configured Session Limit	Current Session Count Displayed by **show services l2tp summary** 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 18 on page 209, 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 18: Scenario 2, Tunnel 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	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 19 on page 210, 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 19: 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 20 on page 210, 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 20: 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

Table 20: Scenario 4, Session-Limit Group Limit *(continued)*

Level	Configured Session Limit	Current Session Count Displayed by **show services l2tp summary** Command	Session Limit Check Result
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

3. Next, suppose you remove client ent1-serviceA from the session-limit group. The group session capacity increases to 5000 sessions:

   6000 - 1000 = 5000

4. 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 21 on page 212, 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 21: 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.

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.

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.

You can use the **show services l2tp summary extensive** command to display the configured sessions limit for a tunnel:

```
user@host> show services l2tp tunnel extensive
```

```
...
    Max sessions: 32000, Window size: 4, Hello interval: 60
...
```

The displayed limit for configured sessions is set to the lowest of the following configured session values:

- Global (chassis)—(LAC and LNS) **set services l2tp tunnel** maximum-sessions ***number***
- Tunnel profile (individual tunnel)—(LAC and LNS) **set access tunnel-profile** ***profile-name*** **tunnel** ***tunnel-id*** max-sessions ***number*****]**
- RADIUS—(LAC and LNS) Value of VSA 26–33, Tunnel-Max-Sessions
- Host profile—(LNS only) **set access profile** ***profile-name*** **client** ***client-name*** **l2tp** maximum-sessions-per-tunnel

The configured values determine the field value starting in the following Junos OS releases: 19.2R3, 19.3R3, 19.4R3, 20.1R2, 20.2R2, and 20.3R1. In earlier releases, the field displays the host profile value for the LNS, but it displays a fixed value of 512,000 for the LAC.

> **NOTE:** After a GRES, a unified ISSU, or a restart of the jl2tpd process, the value of this field is accurate only after a new session comes up on the tunnel. Until that happens, the field shows a value of 65,535 instead of the configured value.
>
> Suppose you have two tunnels, tunnel A and tunnel B. A GRES takes place, and the field for each tunnel shows 65,535. When a new session comes up on tunnel B, the value for that tunnel updates to the configured value. For tunnel A, the field continues to show 65,535 until that tunnel gets a new session.

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
```

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;
    }
}
```

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 220.

2. (Optional) Configure how sessions are load-balanced among tunnels.

See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 221.

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 221.

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
```

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
```

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
```

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.

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.

Release History Table

Release	Description
20.3R1	The configured values determine the field value starting in the following Junos OS releases: 19.2R3, 19.3R3, 19.4R3, 20.1R2, 20.2R2, and 20.3R1.
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

L2TP for Subscriber Access Overview | **153**

Configuring an L2TP LAC | **183**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

Configuring the L2TP Destination Lockout Timeout | **180**

Specifying a Tunnel Profile in a Domain Map

L2TP Session Limits and Load Balancing for Service Interfaces | **283**

Configuring L2TP Tunnel Groups

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | **278**

Domain Mapping Overview

LAC Interoperation with Third-Party LNS Devices | **189**

Filtering RADIUS Attributes and VSAs from RADIUS Messages

L2TP Subscriber Access Lines and Connection Speeds

IN THIS SECTION

- Subscriber Access Line Information Handling by the LAC and LNS Overview | **225**
- Transmission of Tx and Rx Connection Speeds from LAC to LNS | **236**
- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal | **244**
- Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS | **244**
- Configuring the Reporting and Processing of Subscriber Access Line Information | **247**
- Preventing the LAC from Sending Calling Number AVP 22 to the LNS | **252**
- Override the Calling-Station-ID Format for the Calling Number AVP | **253**
- Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds | **255**

Subscriber Access Line Information Handling by the LAC and LNS Overview

IN THIS SECTION

- Access Line Information Forwarding | **226**
- Access Line Information AVPs | **227**
- Connection Speed Updates on the LAC | **234**
- Connection Speed Updates on the LNS | **235**
- Interaction Between Global and Per-Destination Configurations | **236**

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. The access node is typically a DSLAM for DSL access networks or, starting in Junos OS Release 19.3R1, an ONT/ONU for PON access networks. See the following references for more information about DSL Forum VSAs and L2TP AVPs:

- RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*
- RFC 5515, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*
- RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*
- RFC 6320 Draft Extension, *Access Extensions for the Access Node Control Protocol*
- Broadband Forum technical report TR-101, *Migration to Ethernet-Based Broadband Aggregation*

Access Line Information Forwarding

In the network topology shown in Figure 16 on page 227, 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.

> **NOTE:** Starting in Junos OS Release 19.3R1, the access nodes for PON subscriber access lines (such as ONTs and ONUs) are supported in this same scenario, in addition to the previously supported DSL access nodes.

Figure 16: Sample L2TP Network Topology

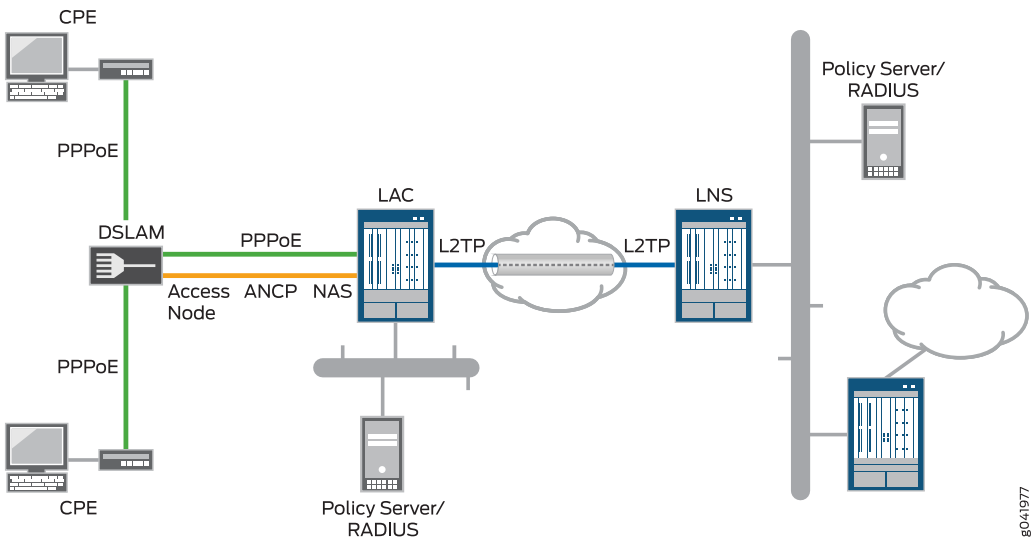

Access Line Information AVPs

L2TP supports the AVPs listed in Table 22 on page 227 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 access node. The content of the ICRQ message generally varies between DSL access lines and PON access lines. AVPs, 1, 2, 3, and 6 are used for access line identification for both DSL and PON. If PON information is reported using DSL AVPs, then the content is the same as it would be for DSL access.

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 22: L2TP AVPs That Provide Subscriber Access Line Information

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
1 Agent-Circuit-Id	Identifier for the subscriber agent circuit ID (ACI) that corresponds to the access node interface from which subscriber requests are initiated. 2-63 octet string	ICRQ DSL, PON	26-3561-1

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
2 Agent-Remote-Id	Unique identifier for the subscriber associated with the access node interface from which requests are initiated. 2-63 octet string	ICRQ DSL, PON	26-3561-2
3 Access-Aggregation-Circuit-ID-ASCII	ASCII identifier for the subscriber access line, based on its network-facing logical appearance If the string begins with a # sign, then the remainder of the string represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. The string is used as the name of a CoS Level 2 interface set that groups subscribers.	ICRQ DSL, PON	26–3561-3
6 Access-Aggregation-Circuit-ID-Binary	Binary identifier for the subscriber access line 32- bit or 64-bit string	ICRQ DSL, PON	26–3561-6
97 Connect-Speed-Update	Data structure listing remote session id and the current transmit and receive connection speeds in bits per second.	CSUN, CSURQ	(none)
98 Connect-Speed-Update-Enable	Value does not matter: presence indicates support for CSUN, CSURQ message types for this session.	ICRQ	(none)

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
129 Actual-Data-Rate-Upstream	Actual upstream data rate of the subscriber's synchronized DSL link, in bps 64-bit unsigned integer; data rate in bits per sec	ICRQ DSL	26-3561-129
130 Actual-Data-Rate-Downstream	Actual downstream data rate of the subscriber's synchronized DSL link, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-130
131 Minimum-Data-Rate-Upstream	Minimum upstream data rate configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-131
132 Minimum-Data-Rate-Downstream	Minimum downstream data rate configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-132
133 Attainable-Data-Rate-Upstream	Upstream data rate that the subscriber can attain, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-133
134 Attainable-Data-Rate-Downstream	Downstream data rate that the subscriber can attain, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-134
135 Maximum-Data-Rate-Upstream	Maximum upstream data rate configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-135

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
136 Maximum-Data-Rate-Downstream	Maximum downstream data rate configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-136
137 Minimum-Data-Rate-Upstream-Low-Power	Minimum upstream data rate in low power state configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-137
138 Minimum-Data-Rate-Downstream-Low-Power	Minimum downstream data rate in low power state configured for the subscriber, in bps 64-bit unsigned integer	ICRQ DSL	26-3561-138
139 Maximum-Interleaving-Delay-Upstream	Maximum one-way upstream interleaving delay configured for the subscriber, in milliseconds 32-bit unsigned integer	ICRQ DSL	26-3561-139
140 Actual-Interleaving-Delay-Upstream	Subscriber's actual one-way upstream interleaving delay, in milliseconds 32-bit unsigned integer	ICRQ DSL	26-3561-140
141 Maximum-Interleaving-Delay-Downstream	Maximum one-way downstream interleaving delay configured for the subscriber, in milliseconds 32-bit unsigned integer	ICRQ DSL	26-3561-141

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
142 Actual-Interleaving-Delay-Downstream	Subscriber's actual one-way downstream interleaving delay, in milliseconds 32-bit unsigned integer	ICRQ DSL	26-3561-142
144 Access-Loop-Encapsulation	Encapsulation used by the subscriber associated with the access node interface from which requests are initiated Three one-octet encodings for data link, encapsulation 1, and encapsulation 2.	ICRQ DSL	26-3561-144
145 ANCP-Access-Line-Type (This corresponds to the ANCP DSL-Type TLV.)	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 DSL	26-3561-145

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
146 PON-Access-Type	Type of PON access line in use: • **0**—OTHER • **1**—GPON • **2**—XG-PON1 • **3**—TWDM-PON • **4**—XGS-PON • **5**—WDM-PON • **7**—UNKNOWN 32-bit unsigned integer	ICRQ PON	26–3561–146
147 ONT/ONU-Average-Data-Rate-Downstream	Average downstream data rate for ONT/ONU, in bps 64-bit unsigned integer	ICRQ PON	26–3561–147
148 ONT/ONU-Peak-Data-Rate-Downstream	Peak downstream data rate for ONT/ONU, in bps 64-bit unsigned integer	ICRQ PON	26–3561–148
149 ONT/ONU-Maximum-Data-Rate-Upstream	Maximum upstream data rate for ONT/ONU, in bps 64-bit unsigned integer	ICRQ PON	26–3561–149
150 ONT/ONU-Assured-Data-Rate-Upstream	Assured upstream data rate for ONT/ONU, in bps 64-bit unsigned integer	ICRQ PON	26–3561–150
151 PON-Tree-Maximum-Data-Rate-Upstream	Maximum upstream data rate for the PON tree, in bps 64-bit unsigned integer	ICRQ PON	26–3561–151

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
152 PON-Tree-Maximum-Data-Rate-Downstream	Maximum downstream data rate for the PON tree, in bps 64-bit unsigned integer	ICRQ PON	26–3561–152
155 Expected-Throughput-Upstream	Expected upstream throughput, which is the net data rate reduced by expected rate loss, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–155
156 Expected-Throughput-Downstream DSL	Expected upstream throughput, which is the net data rate reduced by expected rate loss, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–156
157 Attainable-Expected-Throughput-Upstream	Maximum attainable expected upstream throughput, in Kbps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–157
158 Attainable-Expected-Throughput-Downstream	Maximum attainable expected downstream throughput, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–158
159 Gamma-Data-Rate-Upstream	Actual upstream data rate (net data rate) for the local loop, adjusted down by any throughput capability limitations, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–159

Table 22: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

L2TP AVP Type L2TP AVP Name	Description	L2TP Message Type Access Line Support	Corresponding DSL Forum VSA
160 Gamma-Data-Rate-Downstream	Actual downstream data rate (net data rate) for the local loop, adjusted down by any throughput capability limitations, in Kbps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–160
161 Attainable-Gamma-Data-Rate-Upstream	Maximum attainable upstream data rate (net data rate) for the local loop, adjusted down by any throughput capability limitations, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–161
162 Attainable-Gamma-Data-Rate-Downstream	Maximum attainable downstream data rate (net data rate) for the local loop, adjusted down by any throughput capability limitations, in bps 64-bit unsigned integer	ICRQ DSL (G.fast)	26–3561–162
254 IWF Session	Four-octet field indicating whether or not the internetworking function has been performed for the subscriber's PPPoA over PPPoE session	ICRQ DSL	26-3561–254

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**.

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 244.

- 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 23 on page 240. 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 session scheduler 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 255 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 23 on page 240 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 23: 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	• ancp • none • pppoe-ia-tags • service-profile • static (default)	• 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	• actual (default) • ancp • none • pppoe-ia-tags	• 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	• 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 24 on page 242.

Table 24: 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.
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 25 on page 243.

Table 25: 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 26 on page 243.

Table 26: 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.

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
```

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 236 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 255.

- (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
```

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.

Starting in Junos OS Release 19.3R1, AVPs for PON and G.fast access lines are supported, corresponding to the Broadband Forum PON and G.fast TLVs.

> **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
...
```

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
```

Override the Calling-Station-ID Format for the Calling Number AVP

The LAC sends information about the access line or the subscriber to the LNS in L2TP Calling Number AVP 22. This AVP is conveyed in the incoming-call request (ICRQ) packet when the L2TP session is being established. AVP 22 by default identifies the access node interface that is connected to the customer in the access network; this is the agent circuit identifier or ACI. The LAC receives the ACI in the PPPoE Active Discovery Request (PADR) packet from the L2TP client as DSL Forum Agent-Circuit-ID VSA [26-1].

Alternatively, you can use the **calling-station-id-format** statement to change the values sent in the AVP. For example, you might specify that the agent remote identifier (ARI) received in the PADR as DSL Forum Agent-Remote-ID VSA [26-2] is used instead of the agent circuit identifier, that both are used, or that additional attributes are included. The set of values used in the AVP is known as the Calling-Station-ID format. When this is configured, then the value of the AVP is subsequently sent to the RADIUS server as Calling-Station-ID attribute (31).See *Configuring a Calling-Station-ID with Additional Options* for more information.

In some cases you may want the value of Calling Number AVP 22 to be independent from the RADIUS attribute value. You can do this by overriding the configured Calling-Station-ID format for the value. Use the **remote-circuit-id-format** statement to specify a different format for the AVP: the ACI, the ARI, or both the ACI and ARI from the PADR packet.

You can also configure fallback values that are sent in the Calling Number AVP when the values you configure with the **remote-circuit-id-format** statement are not present in the PADR. You can configure the fallback option to send the configured Calling-Station-ID or the default underlying interface as the calling number AVP.

Before you begin:

- Configure an access profile.
- Configure L2TP.
- Configure RADIUS.

To configure the override in the access profile:

1. Configure the LAC to send the calling number AVP using the configured remote circuit ID format instead of the Calling-Station-ID format.

   > **NOTE:** The **override** statement fails commit check if you have not configured the **remote-circuit-id-format** statement.

   ```
   user@host# set access profile profile-name override calling-station-id remote-circuit-id
   ```

2. Configure the format of the values that override the Calling-Station-ID in AVP 22. You can configure the format to include the ACI, the ARI, or both the ACI and ARI.

   ```
   user@host# set access profile profile-name radius options remote-circuit-id-format agent-circuit-id
   user@host# set access profile profile-name radius options remote-circuit-id-format agent-remote-id
   ```

   Table 27 on page 254 describes the attributes sent in calling number AVP 22 based on the attributes received in the PADR and the format configured in the **remote-circuit-id-format** configuration statement.

   Table 27: Attributes Sent as Calling Number AVP Based on Remote Circuit ID Format and Attributes Received in PADR

Remote Circuit ID Format	Attributes Received in PADR	Attributes Sent in Calling Number AVP
Agent-Circuit-ID	Agent-Circuit-ID, Agent-Remote-ID	Agent-Circuit-ID
Agent-Remote-ID	Agent-Circuit-ID, Agent-Remote-ID	Agent-Remote-ID
Agent-Circuit-ID, Agent-Remote-ID	Agent-Circuit-ID, Agent-Remote-ID	Agent-Circuit-ID, Agent-Remote-ID
Agent-Circuit-ID, Agent-Remote-ID	Agent-Circuit-ID	Agent-Circuit-ID
Agent-Circuit-ID, Agent-Remote-ID	Agent-Remote-ID	Agent-Remote-ID

3. (Optional) Configure the fallback value to be used. Fallback is triggered if the ACI and ARI are not present in the PADR but are configured in the remote circuit ID format. You can configure the LAC to send the configured Calling-Station-ID or the default underlying interface in the Calling number AVP 22 when fallback is triggered.

   ```
   user@host# set access profile profile-name remote-circuit-id-fallback configured-calling-station-id
   ```

```
user@host# set access profile profile-name remote-circuit-id-fallback default
```

The remote circuit ID format determines what triggers the fallback. Table 28 on page 255 shows the fallback trigger based on the remote circuit ID format.

Table 28: Fallback Trigger for Remote Circuit ID Format

Remote Circuit ID Format	Fallback Trigger
Agent-Circuit-ID	Agent-Circuit-ID is empty
Agent-Remote-ID	Agent-Remote-ID is empty
Agent-Circuit-ID, Agent-Remote-ID	Both Agent-Circuit-ID and Agent-Remote-ID are empty

4. (Optional) Configure an alternative delimiter character that the router uses to separate the concatenated values in the resulting remote circuit ID string when more than one value is specified in the remote circuit ID format. The default delimiter is a hash character (#).

   ```
   user@host#set access profile profile-name remote-circuit-id-delimiter "delimiter"
   ```

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 tx-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 Kbps x 1005 = 10,050,000 bps.
  - Tx is the second parameter passed by the VSA multiplied by the configured multiplier:
    15000 Kbps x 1003 = 15,045,000 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 Kbps x 1005 = 10,050,000 bps.
  - Because the VSA does not pass a second parameter, 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, 0)”. The following values are stored in the SDB:
  - Rx is the first parameter passed by the VSA multiplied by the configured multiplier:
    10000 Kbps x 1005 = 10,050,000 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 Kbps x 1005 = 10,050,000 bps.
  - Tx is the second parameter passed by the VSA multiplied by the configured multiplier:
    15000 Kbps x 1003 = 15,045,000 bps.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, AVPs for PON and G.fast access lines are supported, corresponding to the Broadband Forum PON and G.fast TLVs.
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.
17.4R1	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.
17.4R1	Starting in Junos OS Release 17.4R1, RFC 5515 AVP extensions are also supported on the LNS.
17.2R1	Starting in Junos OS Release 17.2R1, the LAC fallback procedure is as described in Table 24 on page 242.
17.2R1	Starting in Junos OS Release 15.1R1, the LAC fallback procedure is as described in Table 25 on page 243.
17.2R1	Starting in Junos OS Release 13.3R1, the LAC fallback procedure is as described in Table 26 on page 243.
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.
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).
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.
13.3R1	Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release, as described in Table 23 on page 240.
13.3R1	Starting in Junos OS Release 13.3R1, availability and support for methods vary by Junos OS Release.

RELATED DOCUMENTATION

- L2TP for Subscriber Access Overview | 153
- Configuring an L2TP LAC | 183
- LAC Tunnel Selection Overview | 191
- *DSL Forum Vendor-Specific Attributes*
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- *RADIUS Servers and Parameters for Subscriber Access*
- *Filtering RADIUS Attributes and VSAs from RADIUS Messages*

L2TP LNS Inline Service Interfaces

IN THIS SECTION

- Configuring an L2TP LNS with Inline Service Interfaces | 261
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266
- Configuring an L2TP Access Profile on the LNS | 268
- Configuring a AAA Local Access Profile on the LNS | 270
- Configuring an Address-Assignment Pool for L2TP LNS with Inline Services | 271
- Configuring the L2TP LNS Peer Interface | 273
- Enabling Inline Service Interfaces | 273
- Configuring an Inline Service Interface for L2TP LNS | 275
- Configuring Options for the LNS Inline Services Logical Interface | 276
- LNS 1:1 Stateful Redundancy Overview | 277
- Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278
- Verifying LNS Aggregated Inline Service Interface 1:1 Redundancy | 280
- L2TP Session Limits and Load Balancing for Service Interfaces | 283
- Example: Configuring an L2TP LNS | 287
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301
- Applying Services to an L2TP Session Without Using RADIUS | 304
- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions | 313
- Configuring a Dynamic Profile for Dynamic LNS Sessions | 314

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 266.

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 263.

3. Configure inline IP reassembly.

   See "Configuring IP Inline Reassembly for L2TP" on page 318.

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 268.

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 270.

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 271.

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 273.

8. Enable inline service interfaces on an MPC.

   See "Enabling Inline Service Interfaces" on page 273.

9. Configure a service interface.

   See "Configuring an Inline Service Interface for L2TP LNS" on page 275.

10. Configure options for each inline service logical interface.

    See "Configuring Options for the LNS Inline Services Logical Interface" on page 276.

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 278

12. Configure the L2TP tunnel group.

    See “Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces” on page 301.

13. (Optional) Configure a dynamic profile that dynamically creates L2TP logical interfaces.

    See “Configuring a Dynamic Profile for Dynamic LNS Sessions” on page 314.

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 313.

15. (Optional) Specify how many times L2TP retransmits unacknowledged control messages.

    See “Configuring Retransmission Attributes for L2TP Control Messages” on page 160.

16. (Optional) Specify how long a tunnel can remain idle before being torn down.

    See “Setting the L2TP Tunnel Idle Timeout” on page 178.

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 177.

18. (Optional) Specify how long the L2TP retains information about terminated dynamic tunnels, sessions, and destinations.

    See “Setting the L2TP Destruct Timeout” on page 178.

19. (Optional) Configure the L2TP destination lockout timeout.

    See “Configuring the L2TP Destination Lockout Timeout” on page 180.

20. (Optional) Configure L2TP tunnel switching.

    See “Configuring L2TP Tunnel Switching” on page 175.

21. (Optional) Prevent the creation of new sessions, destinations, or tunnels for L2TP.

    See “Configuring L2TP Drain” on page 181.

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 321.

23. (Optional) Enable SNMP statistics counters.

    See “Enabling Tunnel and Global Counters for SNMP Statistics Collection” on page 163.

24. (Optional) Configure trace options for troubleshooting the configuration.

    See “Tracing L2TP Events for Troubleshooting” on page 323.

You also need to configure CoS for LNS sessions. For more information, see *Configuring Dynamic CoS for an L2TP LNS Inline Service*.

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
   ```

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 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.

5. 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
   ```

6. (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:

1. Specify the logical inline service interface.

   ```
   [edit interfaces si-slot/pic/port]
   user@host# edit unit logical-unit-number
   ```

2. 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
   ```

3. 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 266 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.

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 263 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.

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.

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
```

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.

For more information about address assignment pools, see *Address-Assignment Pools Overview* and *Address-Assignment Pool Configuration Overview*.

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
```

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.

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.

Although the range of bandwidth values is 1 Gbps through 400 Gbps, you cannot configure the bandwidth in absolute numbers such as 12,345,878,000 bps. You must use the options available in the CLI statement:

- **1g**
- **10g** through **100g** in 10 Gbps increments: **10g**, **20g**, **30g**, **40g**, **50g**, **60g**, **70g**, **80g**, **90g**, **100g**
- **100g** through **400g** in 100 Gbps increments: **100g**, **200g**, **300g**, **400g**

The maximum bandwidth available varies among MPCs, as shown in Table 29 on page 274. A system log message is generated when you configure a bandwidth higher than is supported on the MPC.

Table 29: Maximum Bandwidth for Inline Services per MPC

MPC	Maximum Supported Bandwidth
MPC2E NG, MPC2E NG Q,	40 Gbps
MPC3E NG, MPC3E NG Q	40 Gbps
100GE and 40GE MPC3 and MICs	40 Gbps
MPC4E	40 Gbps
MPC5E	40 Gbps
MPC6E	40 Gbps
MPC7E	240 Gbps
MPC8E	240 Gbps 400 Gbps in 1.6 Tbps upgraded mode
MPC9E	400 Gbps

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 optionally specify the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services. Starting in Junos OS Release 16.2, you are not required to explicitly specify a bandwidth for L2TP LNS tunnel traffic using inline services. When you do not specify a bandwidth, the maximum bandwidth supported on the PIC is automatically available for the inline services; inline services can use up to this maximum value. In earlier releases, you must specify a bandwidth when you enable inline services with the **inline-services** statement.

```
[edit chassis fpc slot-number pic number]
user@host# set inline-services bandwidth bandwidth-value
```

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
```

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 314.

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
```

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, asi*X*.*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.

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, asi*X*.*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 asi*X*.*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 273 and "Configuring an Inline Service Interface for L2TP LNS" on page 275.
- 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
```

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: Sendbcast-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
```

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.

Example: Configuring an L2TP LNS

IN THIS SECTION

- Requirements | 287
- Overview | 288
- Configuration | 290

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

This L2TP LNS example requires the following hardware and software:

- MX Series 5G Universal Routing Platform
- 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 30 on page 287 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 30: 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

Table 30: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example *(continued)*

VSA Name [Number]	Order	Value
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-lns-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-encapuslation** 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.

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) Specify a dynamic 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
```

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-lns-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
```

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
   ```

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 User Guide for information about using variables for IPv6-only and dual-stack addressing in dynamic profiles.

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.
16.2R1	Starting in Junos OS Release 16.2, you are not required to explicitly specify a bandwidth for L2TP LNS tunnel traffic using inline services.

RELATED DOCUMENTATION

L2TP for Subscriber Access Overview | **153**

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses | **144**

L2TP Session Limits Overview | **207**

Local and Remote Service Activation and Deactivation Using the CLI

Junos OS Feature Licenses

Session Options for Subscriber Access

IP Packet Reassembly on Inline Service Interfaces

IN THIS SECTION

- IP Packet Fragment Reassembly for L2TP Overview | 317
- Configuring IP Inline Reassembly for L2TP | 318

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 reassembly of 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.

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

Configuring an L2TP LNS with Inline Service Interfaces | 261

Protocols and Applications Supported on the MPC1E for MX Series Routers

Peer Resynchronization After an L2TP Failover

IN THIS SECTION

- L2TP Failover and Peer Resynchronization | 320
- Configuring the L2TP Peer Resynchronization Method | 321

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.

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 tunnel]
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 tunnel]
user@host# set failover-resync silent-failover
```

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.
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 for Subscriber Access Overview | 153

Tracing L2TP Events for Troubleshooting

IN THIS SECTION

- Configuring the L2TP Trace Log Filename | 324
- Configuring the Number and Size of L2TP Log Files | 325
- Configuring Access to the L2TP Log File | 325
- Configuring a Regular Expression for L2TP Messages to Be Logged | 326

- Configuring Subscriber Filtering for L2TP Trace Operations | 326
- Configuring the L2TP Tracing Flags | 328
- Configuring the Severity Level to Filter Which L2TP Messages Are Logged | 328

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.

The following topics describe how to configure all aspects of tracing L2TP operations:

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
```

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
```

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
```

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
```

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
```

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
```

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
```

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

L2TP for Subscriber Access Overview | 153

5
CHAPTER

Configuring MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

IN THIS SECTION

- Pseudowire Subscriber Logical Interfaces Overview | 331
- Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview | 334
- Configuring a Pseudowire Subscriber Logical Interface | 338
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router | 340
- Configuring a Pseudowire Subscriber Logical Interface Device | 341
- Changing the Anchor Point for a Pseudowire Subscriber Logical Interface Device | 344
- Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface | 346
- Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces | 347
- Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces | 348
- Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface | 350

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 17 on page 332 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 17: MPLS Access Network with Subscriber Management Support

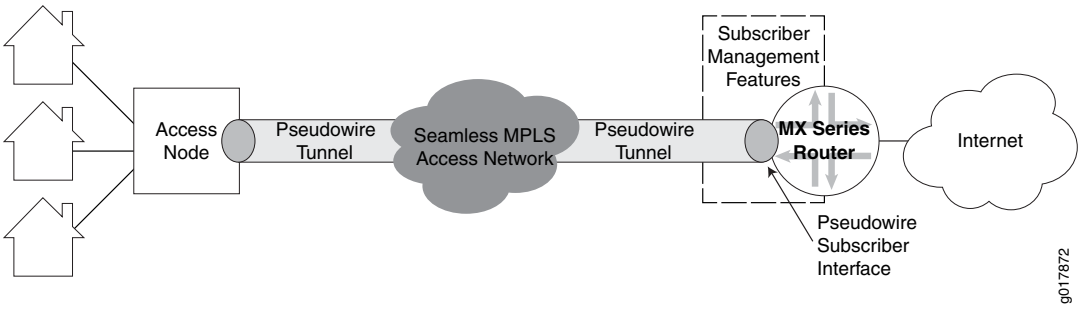

Figure 18 on page 333 shows the protocol stack for a pseudowire subscriber logical interface. The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD), and supports a circuit-oriented Layer 2 protocol (either Layer 2 VPN or Layer 2 circuit). The Layer 2 protocol provides the transport and service logical interfaces, and supports the protocol family (IPv4, IPv6, or PPPoE).

Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the support for pseudowire subscriber service interface over redundant logical tunnels is introduced in Layer 3 VPNs and draft-rosen multicast VPNs. Earlier, Layer 3 VPNs provided support for pseudowire subscriber services over logical tunnel interfaces only, and these interfaces used unicast routing protocols, such as OSPF or BGP. With this support, you can provision a multicast routing protocol, Protocol Independent Multicast (PIM), on the pseudowire subscriber interfaces, which gets terminated on the virtual routing and forwarding (VRF) routing instance. Additionally, there is an increase in the scaling numbers of the pseudowire logical interface devices that provides additional resiliency support for pseudowire subscriber interfaces on redundant logical tunnel interfaces.

NOTE: When a pseudowire subscriber service interface is anchored to a redundant logical tunnel whose member interface (or FPC) does not exist, the tunnel interface comes down. In such cases, the pseudowire interfaces (physical and logical) should also be down, but however, the pseudowire subscriber logical interface state remains up, although the Layer 2 circuit services, such as ping toward a customer edge (CE) device from the service side of the pseudowire subscriber service interface, are not available.

This is because the transport side of the pseudowire subscriber logical interface stays up causing the services to be up.

Figure 18: Pseudowire Subscriber Interface Protocol Stack

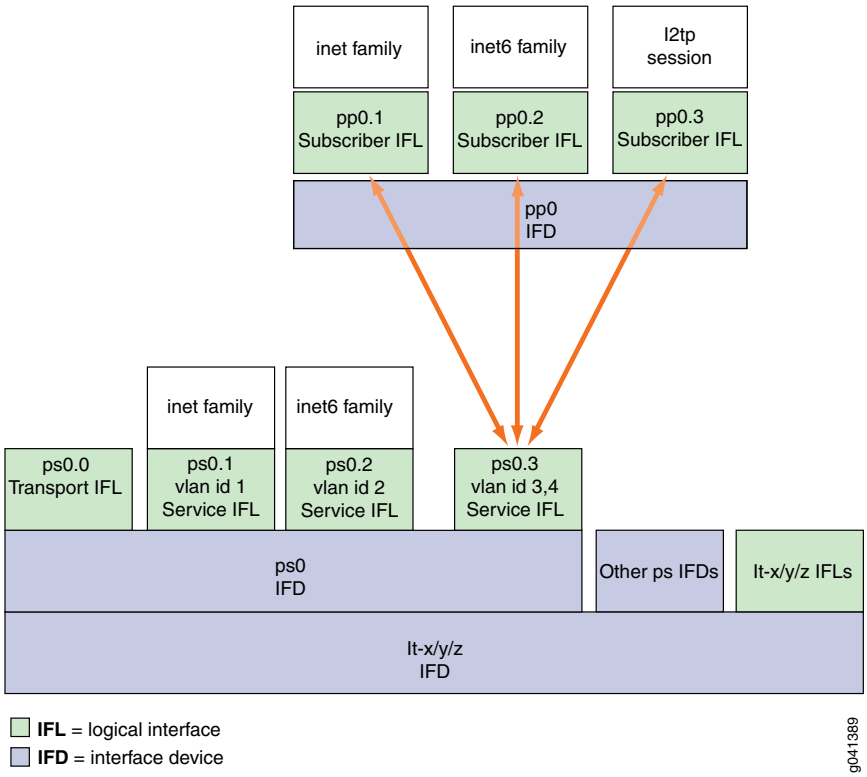

The pseudowire configuration is transparent to the subscriber management applications and has no impact on the packet payloads that are used for subscriber management. Subscriber applications such as DHCP and PPPoE can be stacked over Layer 2 similar to the way in which they are stacked over a physical interface.

Starting with Junos OS release 16.1R1, **family inet** and **family inet6** are supported on the services side of an MPLS pseudowire subscriber as well as non-subscriber logical interface.

Starting with Junos OS Release 16.1R1, Inline IPFIX is supported on the services side of an MPLS pseudowire subscriber logical interface.

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.

Prior to Junos OS Release 19.1R1, the only supported encapsulation type on the pseudowire subscriber interfaces included:

- Transport logical interfaces—Circuit cross-connect (CCC) encapsulation.
- **Service logical interfaces**:
  - Ethernet VPLS encapsulation
  - VLAN bridge encapsulation

- VLAN VPLS encapsulation

Starting in Junos OS Release 19.1R1, additional encapsulations are added to the pseudowire subscriber transport and service logical interfaces. The transport logical interface supports Ethernet VPLS encapsulation, and provisions for terminating the interface on the **l2backhaul-vpn** routing-instance. The service logical interface supports circuit cross-connect (CCC) encapsulation, and provisions for terminating the interface on locally switched Layer 2 circuits.

With the support of additional encapsulation types, you can benefit from demux of a **l2backhaul** VPN into multiple VPN services, such as Layer 2 circuit and Layer 3 VPN. Because pseudowire subscriber interfaces are anchored on redundant logical tunnels, this enhancement also provides line card redundancy.

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.

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.

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.

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.

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.

Prior to Junos OS Release 18.3R1, you could specify a maximum of 2048 pseudowire subscriber redundant logical tunnel interface devices for an MX Series router. Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the pseudowire redundant logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.

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 underlaying 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 19 on page 336 provides the details of pseudowire subscriber logical interface stacking over a redundant logical tunnel interface.

Figure 19: Pseudowire Subscriber Logical Interface Stacking over Redundant Logical Tunnel Interface

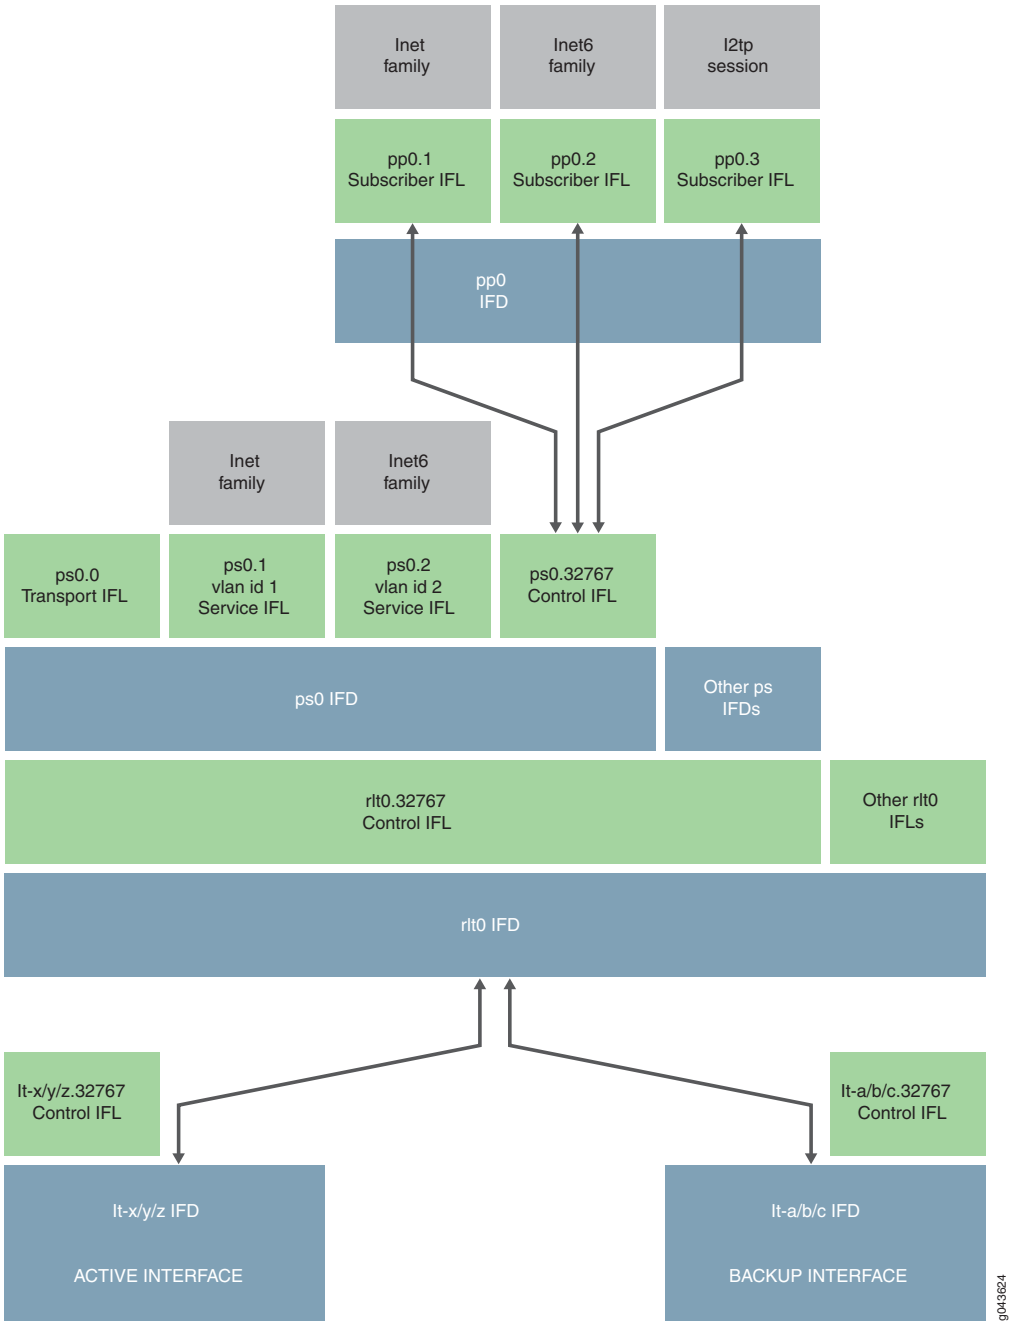

NOTE: Static service ifl is not stacked over transport ifl when RLT is used.

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 rlt*X* 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 an 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.

NOTE: You cannot disable the underlying redundant logical tunnel (rlt) interface or the underlying logical tunnel (lt) interface when a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

Starting in Junos OS Release 18.4R1, the support for inline distribution of single-hop Bidirectional Forwarding Detection (BFD) sessions is extended to pseudowire subscriber over redundant logical tunnel interfaces. For pseudowire subscriber over logical tunnel interfaces, the interfaces are anchored on a single Flexible PIC Concentrator (FPC), as a result, the inline distribution of single-hop BFD sessions is supported by default. With pseudowire redundant logical interfaces, the member logical tunnel interfaces can be hosted on different linecards. Because the distribution address is not available for the redundant logical interfaces, the distribution of single-hop BFD sessions was operated in a centralized mode before Junos OS Release 18.4R1.

With the support for inline distribution of single-hop BFD sessions over pseudowire redundant logical interfaces, there is a scaling advantage of up to 2000 single-hop BFD sessions at an interval of one second, and improvement in detection time enhancing the performance of the sessions.

The BFD operation for pseudowire subscriber over redundant logical interfaces is as follows:

1. When a new BFD session gets added it can either be anchored on an active or a backup FPC.

2. When either of the FPCs fail or reboot, all the sessions hosted on that FPC go down, and re-anchoring is triggered for the next available distribution address. The BFD sessions come back up after the sessions are installed on the other FPC and BFD packet exchange is started.

   However, it is also possible that the sessions on the backup FPC might not go down when active FPC fails depending on the BFD detection time configured, as the forwarding state for the new active FPC might take some time to be programmed.

3. When the active FPC fails, all the BFD sessions get anchored on the backup FPC. Similarly, if the backup FPC fails, all the BFD sessions get anchored on the active FPC.

4. The BFD session re-anchoring is not triggered when the active FPC is online again.

5. With the non-revertive behavior enabled, when the previously active FPC is online again, the sessions are not expected to go down. With the default revertive behavior, it is possible that forwarding state needs to be updated and depending on the detection time configuration, the session may or may not flap.

> **NOTE:**
> Take the following into consideration with the support of inline distribution of single-hop BFD sessions on pseudowire subscriber over logical tunnel interfaces:
>
> - On FPC type MPC 7e, with the activation of 7000 routing instance, it takes about six minutes for the 7000 BGP sessions to get established on the pseudowire subscriber interfaces anchored on redundant logical tunnel interfaces.
> - A new system log error message - **JTASK_SCHED_SLIP** - is recorded during nonstop active routing (NSR). This is expected behavior of NSR with high scale and can be safely ignored, unless there are other issues, such as session flaps, that require action to be taken.

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 340.

2. Configure the pseudowire subscriber logical interface device.

   See "Configuring a Pseudowire Subscriber Logical Interface Device" on page 341.

3. Configure the transport logical interface.

   See "Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface" on page 346.

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 347.
   - To configure Layer 2 VPN signaling, see "Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces" on page 348.

5. Configure the service logical interface.

   See "Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface" on page 350.

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* .

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.

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. Setting the maximum number also defines the interface names for the pseudowire interfaces. When you subsequently configure the interfaces, you must specify the interface names in the range from ps0 up to ps*(device-count - 1)*.

For example, if you set the maximum number of devices to 5, then you can configure only interfaces ps0, ps1, ps2, ps3, and ps4.

Before Junos OS Release 17.2R1, you could specify a maximum of 2048 pseudowire logical interface devices for an MX Series router. Starting in Junos OS Release 17.2R1, on MX Series routers with MPC and MIC interfaces, the pseudowire logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.

Similarly, before Junos OS Release 18.3R1, you could specify a maximum of 2048 pseudowire subscriber redundant logical tunnel (rlt) interface devices for an MX Series router. Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the pseudowire redundant logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.

The PFE hosting the maximum pseudowire logical interface devices 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
```

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 use redundant logical tunnels to provide redundancy for member logical tunnels. You can configure additional optional parameters for the interface device, such as VLAN tagging method, MTU, and gratuitous ARP support.

NOTE: You must create a logical tunnel for the pseudowire logical interface device. If you are using redundant logical tunnels, you must create the redundant tunnel.

To configure the pseudowire subscriber interface device:

1. Specify that you want to configure the pseudowire subscriber logical interface device.

NOTE: The available interface names are determined by the **[edit chassis pseudowire-service device-count]** statement. The names you specify must be in the range ps0 through ps*(device-count - 1)*. If you specify an interface name outside that range, the pseudowire interface is not created.

```
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 or a member link in a redundant logical tunnel. You use the command, **set chassis fpc *slot-number* pic *pic-number* tunnel-services bandwidth *bandwidth*** to enable tunnel services.

NOTE: You cannot disable the underlying logical tunnel (lt) interface or redundant logical tunnel (rlt) interface when a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

```
[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 before 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 encapsulation type for the pseudowire logical interface device.

   Starting in Junos OS Release 19.1R1, you can configure additional encapsulations – Ethernet VPLS and circuit cross-connect-based encapsulations – for the transport and service pseudowire subscriber logical interface devices, respectively.

   ```
   [edit interfaces]
   user@host# set logical-interface-unit encapsulation encapsulation-type
   ```

6. (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.

7. (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.

8. (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 *Understanding Unicast RPF (Routers)* for additional information.

9. Configure additional optional parameters for the pseudowire logical interface device, such as *description*, *apply-groups*, *apply-groups-except*, and traceoptions.

Changing the Anchor Point for a Pseudowire Subscriber Logical Interface Device

You cannot dynamically change an anchor point that has active pseudowire devices stacked above it. You must commit certain changes before you move the anchor point. Examples of this situation include moving the anchor point from one logical tunnel to another logical tunnel, from a logical tunnel to a redundant logical tunnel, and from a redundant logical tunnel to a logical tunnel.

To move the anchor point between logical tunnel interfaces:

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 to the new logical tunnel interface and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point lt-fpc/pic/port
user@host# commit
```

3. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

To move the anchor point from a logical tunnel interface to a redundant logical tunnel interface:

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. Add the new redundant logical tunnel interface and commit.

   a. Create the tunnel and set the maximum number of devices allowed.

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-logical-tunnel device-count count
```

b. Bind each member logical tunnel to the redundant logical tunnel.

> **NOTE:** Redundant logical tunnels require members to be in active-backup mode. The backup logical tunnel must be on a different FPC than the active logical tunnel. For example, if the active tunnel is on FPC 3, then the backup tunnel must be on a different FPC, such as FPC 4.

```
[edit interfaces rltnumber]
user@host# set redundancy-group member-interface lt-fpc/pic/port active
user@host# set redundancy-group member-interface lt-fpc/pic/port backup
```

c. Commit your changes.

```
[edit interfaces rltnumber]
user@host# commit
```

3. Change the anchor on the deactivated pseudowire to the new redundant logical tunnel interface and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point rltnumber
user@host# commit
```

4. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

To move the anchor point from a redundant logical tunnel interface to a logical tunnel interface that is a member of the redundant logical tunnel:

1. Deactivate the stacked pseudowires; this may require bringing down any subscribers using the pseudowires. Delete the redundant logical tunnel interface and commit your changes.

```
[edit interfaces]
user@host# deactivate psnumber
user@host# delete rltnumber
user@host# commit
```

2. Change the anchor on the deactivated pseudowire to the new logical tunnel interface and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point lt-fpc/pic/port
user@host# commit
```

3. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

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. (Optional) Specify the encapsulation method for the transport logical interface.

   Starting in Junos OS Release 19.1R1, you can configure Ethernet VPLS encapsulation, in addition to circuit cross-connect-based encapsulations for pseudowire subscriber transport logical interfaces.

```
[edit interfaces ps0 unit 0]
user@host# set encapsulation ethernet-ccc
user@host# set encapsulation ethernet-vpls
user@host# set family vpls core-facing
```

4. (Optional) Configure the termination of the transport logical interface on **l2backhaul-vpn** routing-instance. This support is enabled from Junos OS Release 19.1R1.

```
[edit routing-instances routing-instance-name]
user@host# set vlan-model one-to-one instance-role access instance-type l2backhaul-vpn interface ps1.0s
user@host# set no-local-switching
```

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
```

For more information about Layer 2 circuits, see *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
```

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 331, 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. (Optional) Specify the encapsulation type for the service logical interface.

   Starting in Junos OS Release 19.1R1, you can configure circuit cross-connect-based encapsulations, in addition to the Ethernet VPLS, VLAN bridge, and VLAN VPLS encapsulations for pseudowire subscriber service logical interfaces.

The pseudowire subscriber service logical interfaces support single-tagged traffic, double-tagged traffic, and list of VLANs on the single logical interface.

```
[edit interfaces ps0]
user@host# set unit 1 encapsulation vlan-ccc
user@host# ser vlan-id vlan-ID
user@host# set vlan-tags outer outer-tag inner inner-tag
user@host# set vlan-id-list vlan-id-list
user@host# set family ccc
```

4. (Optional) Configure filters and policers on the family circuit cross-connect encapsulation.

```
[edit interfaces ps0]
user@host# set unit 1 family ccc filter group
user@host# set unit 1 family ccc filter input input-list
user@host# set unit 1 family ccc filter output output-list
user@host# set unit 1 family ccc policer input
user@host# set unit 1 family ccc policer output
```

5. Configure the VLAN tag IDs.

```
[edit interfaces ps0 unit 1]
user@host# set vlan-tags outer 1 inner 1
```

6. 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
```

7. 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
```

8. (Optional) Configure the termination of the service logical interface on locally switched Layer 2 circuits. This support is enabled from Junos OS Release 19.1R1.

```
[edit protocols]
user@host# set l2circuit local-switching interface ps0.1 encapsulation-type ethernet-vlan
   ignore-encapsulation-mismatch ignore-mtu-mismatch
```

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, additional encapsulations are added to the pseudowire subscriber transport and service logical interfaces. The transport logical interface supports Ethernet VPLS encapsulation, and provisions for terminating the interface on the **l2backhaul-vpn** routing-instance. The service logical interface supports circuit cross-connect (CCC) encapsulation, and provisions for terminating the interface on locally switched Layer 2 circuits.
19.1R1	Starting in Junos OS Release 19.1R1, you can configure additional encapsulations – Ethernet VPLS and circuit cross-connect-based encapsulations – for the transport and service pseudowire subscriber logical interface devices, respectively.
19.1R1	Starting in Junos OS Release 19.1R1, you can configure Ethernet VPLS encapsulation, in addition to circuit cross-connect-based encapsulations for pseudowire subscriber transport logical interfaces.
19.1R1	Starting in Junos OS Release 19.1R1, you can configure circuit cross-connect-based encapsulations, in addition to the Ethernet VPLS, VLAN bridge, and VLAN VPLS encapsulations for pseudowire subscriber service logical interfaces.
18.4R1	Starting in Junos OS Release 18.4R1, the support for inline distribution of single-hop Bidirectional Forwarding Detection (BFD) sessions is extended to pseudowire subscriber over redundant logical tunnel interfaces.
18.3R1	Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the support for pseudowire subscriber service interface over redundant logical tunnels is introduced in Layer 3 VPNs and draft-rosen multicast VPNs.
18.3R1	Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the pseudowire redundant logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.
18.3R1	Starting in Junos OS Release 18.3R1, on MX Series routers with MPC and MIC interfaces, the pseudowire redundant logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.
17.3R1	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.
17.2R1	Starting in Junos OS Release 17.2R1, on MX Series routers with MPC and MIC interfaces, the pseudowire logical interface devices scaling numbers has increased to 7000 devices to provide additional resiliency support.

16.1R1	Starting with Junos OS release 16.1R1, **family inet** and **family inet6** are supported on the services side of an MPLS pseudowire subscriber as well as non-subscriber logical interface.
16.1R1	Starting with Junos OS Release 16.1R1, Inline IPFIX is supported on the services side of an MPLS pseudowire subscriber logical interface.
15.1R3	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.
15.1R3	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.
15.1R3	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.
15.1R3	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.

RELATED DOCUMENTATION

- *Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview*
- *CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces*
- *CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces*
- *Tunnel Interface Configuration on MX Series Routers Overview*
- *Router Chassis Configuration Statements*

CHAPTER 6

Wi-Fi Access Gateways

Wi-Fi Access Gateways

IN THIS SECTION

- Wi-Fi Access Gateway Overview | **358**
- Wi-Fi Access Gateway Deployment Model Overview | **360**
- Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway | **361**
- Wi-Fi Access Gateway Configuration Overview | **362**
- Configuring a Pseudowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway | **363**
- Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation | **365**
- Configuring VLAN Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways | **367**
- Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways | **371**

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.

Starting in Junos OS Release 17.2R1, 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. Figure 20 on page 358 shows a sample topology.

Figure 20: MX Series Router Deployed as a WAG

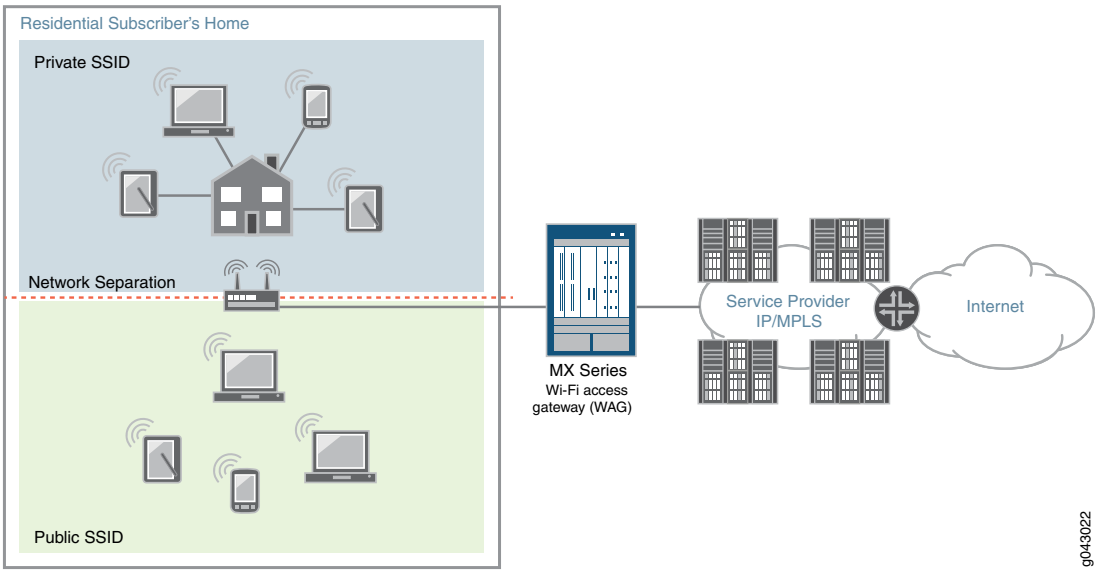

After a WAG has been deployed, service providers can configure the WAG to create secure wireless home network connections for computers, laptops, and other Wi-Fi electronic products (such as game systems, tablets, or mobile phones). WAG offers wireline and mobile service providers the following deployments 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 place 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.

Wi-Fi Access Gateway Deployment Model Overview

Figure 21 on page 360 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 21: MX Series as Wi-Fi Access Gateway Deployment Model

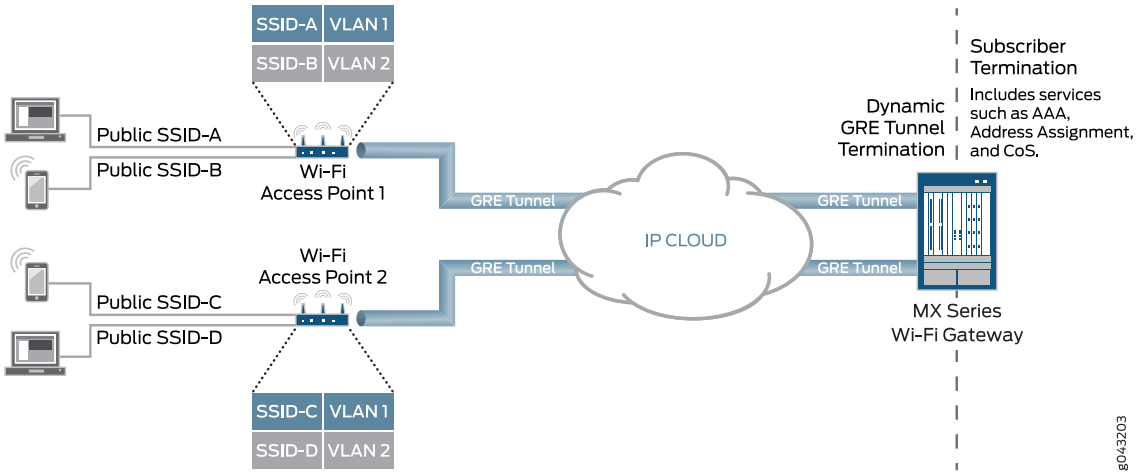

To support the MX Series BNG deployed as a WAG, dynamic-bridged generic routing encapsulation (GRE) tunnels are created and terminated at the BNG when it receives GRE traffic from the wireless access point (WAP). Dynamic Host Configuration Protocol (DHCP) subscribers are transported through GRE tunnels 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.

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway

IN THIS SECTION

- Dynamic VLAN-Tagged Subscribers | **362**
- Untagged Subscribers | **362**

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 (ps*n* IFDs) require that multiple tunnels share the same ps*n* IFD. The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD).

NOTE: The ps*n* 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

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 (ps*n* IFL) (associated with a VLAN ID) on a ps*n* IFD represents multiple GRE tunnels.

A dynamic VLAN demux interface (demux0.*yyyyyyyy*) is created for each VLAN tag and is stacked over the tunnel ps*n* interface (ps*n.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 ps*n* logical interface (ps*n.xxxxxxxx*). There can be multiple subscribers over the same GRE tunnel.

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 Pseudowire Subscriber Logical Interface Device for the Wi-Fi Access Gateway" on page 363.

2. Configure the conditions for enabling dynamic-bridged GRE tunnels.

   See "Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation" on page 365.

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 367.
   - If the subscriber traffic is untagged, see "Configuring Untagged Subscriber Interfaces for Dynamic-Bridged GRE Tunnels on Wi-Fi Access Gateways" on page 371.

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 340.
- 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
```

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 363.
- 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 192.0.2.0/24
```

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.

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
```

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
```

Release History Table

Release	Description
17.2R1	Starting in Junos OS Release 17.2R1, 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.

CHAPTER 7

Fixed Wireless Access Networks

Fixed Wireless Access Networks | 374

Tracing Fixed Wireless Access Events for Troubleshooting | 388

Fixed Wireless Access Networks

IN THIS SECTION

- Fixed Wireless Access Network Overview | 374
- How to Configure Fixed Wireless Access | 383
- Verifying and Monitoring Fixed Wireless Access | 387

Fixed Wireless Access Network Overview

Service providers can manage subscribers over a wireless network to the home instead of having to run fiber to the building. The subscribers are in a fixed location, typically a residence, with customer premises equipment (CPE) that exchanges wireless radio signals with the provider network. The wireless network uses a fiber backhaul tower to handle traffic for the last miles from a hard-wired network to the subscribers. Multiple towers can relay traffic between each other to and from the fiber optic network. Starting in Junos OS Release 19.2R1, MX Series routers acting as BNGs can support subscribers in Third-Generation Partnership Project (3GPP) fixed wireless access networks, enabling the integration of fixed wireless subscribers with the wireline subscriber management backend.

You can find a helpful summary of related terminology in “Terminology for a 3GPP Fixed Wireless Access Network” on page 379.

Figure 22 on page 375 shows a representative topology for a fixed wireless access network.

Figure 22: Fixed Wireless Access Network Topology

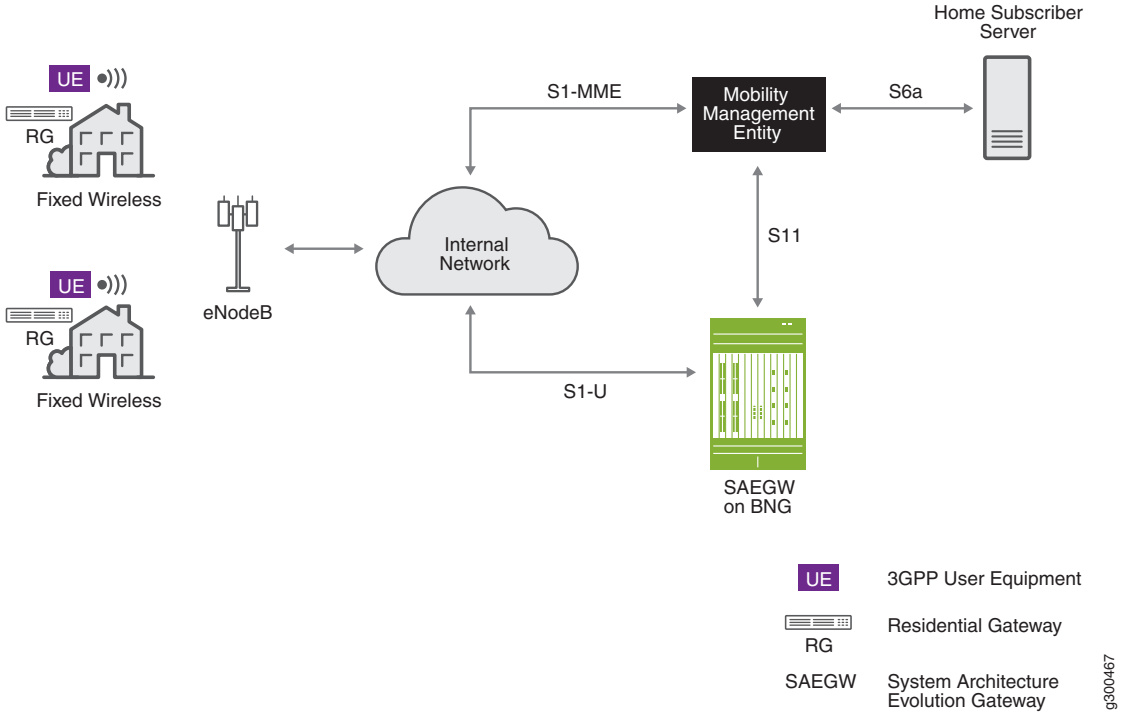

Figure 23 on page 376 shows the GPRS tunneling protocol (GTP) signaling messages in the fixed wireless access network that are necessary to activate and deactivate a default bearer. The messages are actually request and response pairs, but for simplicity the pairs are represented in the figure with bidirectional arrows. The Default Bearer Active line represents the activated default bearer that will carry data traffic.

NOTE: In this implementation, the UE roaming is not supported.

Figure 23: Signaling Messages for Bringing Default Bearer Up and Down

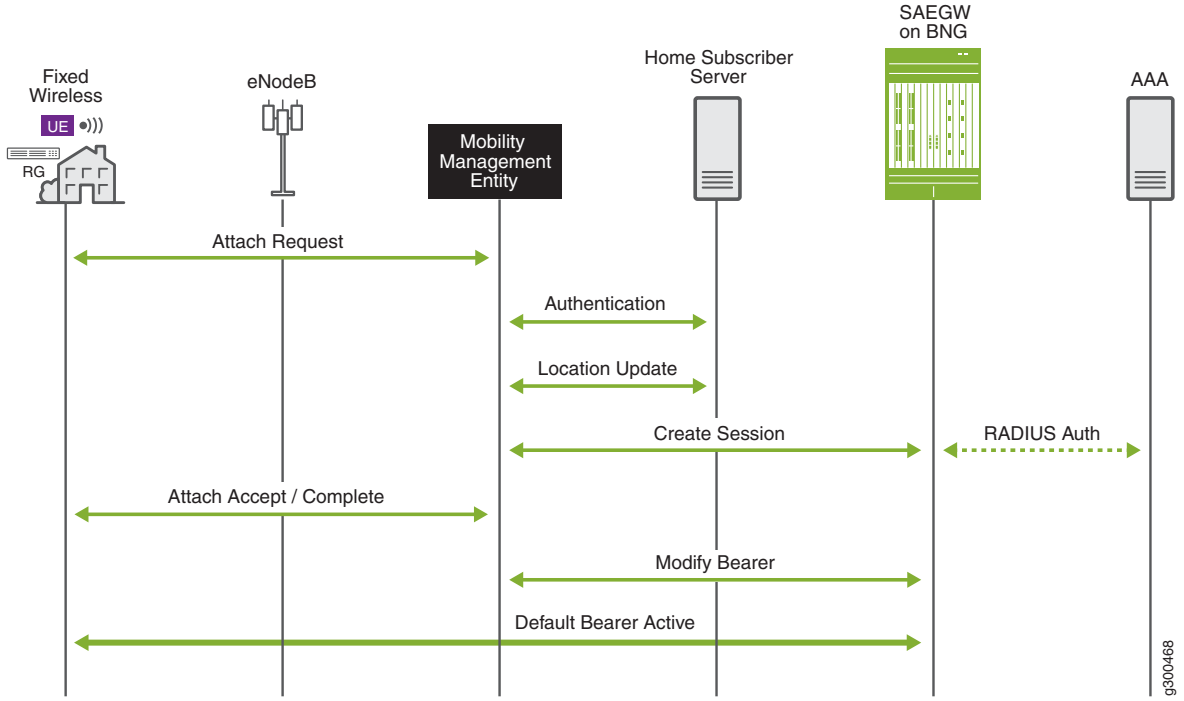

Briefly, the message exchange sequence is the following:

1. The user equipment (UE) sends an Attach Request to the eNodeB, which forwards the message to the mobility management entity (MME). The request includes the APN.

2. The MME exchanges Identity request/response messages with the UE to determine the International Mobile Subscriber Identity (IMSI) or other identifier of the UE.

3. The MME sends the identifier to the Home Subscriber Server (HSS) to authenticate the UE.

4. The MME exchanges Location Update request/response messages with the HSS. In this exchange, the MME provides the HSS with its own address. The HSS acknowledges the update and sends subscription information for the UE from its database.

5. The MME sends a Create Session Request to the System Architecture Evolution gateway (SAEGW). The MME first compares the APN provided by the UE with the APN for which it is authorized according to the subscription information from the HSS. If there is a match, the MME includes that APN in the Create Session Request. If it does not match, then the MME instead includes the APN authorized by the HSS.

6. The SAEGW performs the following actions:

- Validates information elements received in the request.
- Validates the APN requested by the subscriber. If you have configured authentication, the BNG communicates with the RADIUS server to accomplish that task.
- Receives the R-TEID-C, which consists of the IP address of the S11 interface on the MME and the MME's allocated identifier.
- Allocates the local TEIDs:
  - The L-TEID-C is the IP address of the S11 interface on the BNG and the BNG's allocated identifier.
  - The L-TEID-U is the IP address of the S1-U interface on the BNG and the BNG's allocated identifier.
- Allocates an IP address for the pseudowire interface it creates for the session. The address comes from either a locally configured address pool on the BNG or from the RADIUS server.
- Creates the session and sends the Create Session response to the MME. The response includes the IP address that the SAEGW has assigned for the bearer and both L-TEID-C and L-TEID-U.

7. The MME exchanges messages with the eNodeB, which then establishes the bearer component from the UE to eNodeB.

8. The UE signals to the eNodeB that attachment is complete; the eNodeB notifies the MME.

9. The MME and SAEGW exchange Modify Bearer request/response messages to determine the final parameters for the bearer.

   When the BNG receives the request, it creates the dynamic pseudowire (**ps**) interface that will receive the GTP-U encapsulated data packets from the eNodeB. The BNG creates one dynamic **ps** interface per UE. The BNG also receives the R-TEID-U from the eNodeB in the Modify Bearer request. After creating the interface, the BNG sends a Modify Bearer response to the MME.

10. The default bearer is now active and subscriber data traffic can pass back and forth between the UE through eNodeB to the SAEGW and then the connected PDN.

    However, if the residential gateway is a DHCP client, it first begins the DHCP message exchange for the subscriber. The exchange takes place on the default bearer. After DHCP operations have completed to bind the subscriber and provide the DHCP configuration, then data traffic passes over the bearer.

The MME and SAEGW also exchange Delete Session request and response messages. When the SAEGW receives the request, it initiates the subscriber logout process, much as it would for a wireline DHCP subscriber.

Figure 24 on page 378 shows the DHCP connection for the case where the RG creates DHCP subscribers. The DHCP control packets are exceptioned to the Routing Engine as for any other DHCP deployment. The behavior for creating and controlling the DHCP subscribers is the same as for a wireline broadband network. When the subscriber is bound, the UE can then start sending data traffic for the subscriber.

Figure 24: DHCP Subscribers in a Fixed Wireless Access Network

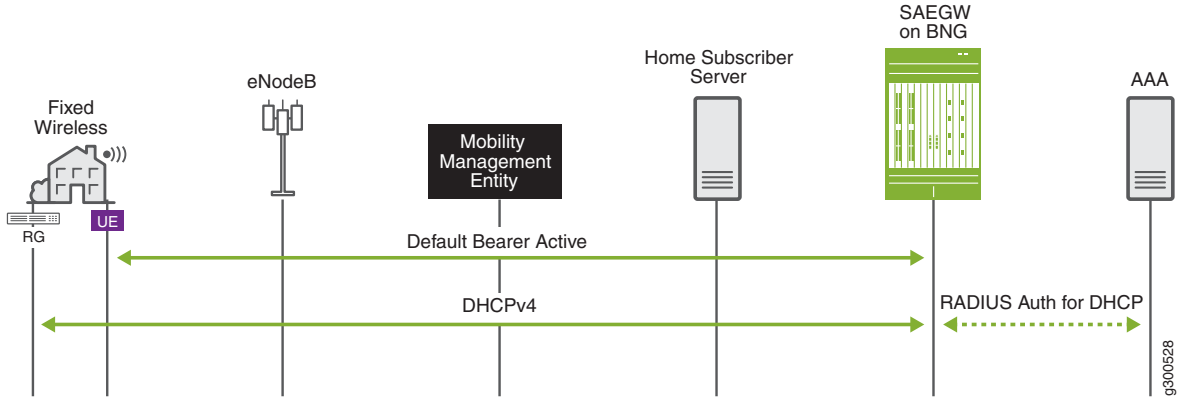

References

For detailed information about all aspects of a fixed wireless network, read the 3GPP technical specifications that define everything. Table 31 on page 378 lists the most relevant specifications.

Table 31: 3GPP Technical Specifications for Fixed Wireless Access Networks

Specification Number	Title
3GPP TS 23.002 (Release 15)	Network architecture
3GPP TS 23.401 (Release 15)	General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
3GPP TS 29.274 (Release 15)	3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C) Stage 3
3GPP TS 29.281 (Release 15)	General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)

3GPP Fixed Wireless Access Terminology

Table 32 on page 379 explains terminology used for 3GPP fixed wireless access networks.

Table 32: Terminology for a 3GPP Fixed Wireless Access Network

Term	Description
3GPP	The 3rd Generation Partnership Project is an international standards organization that develops specifications and protocols for wireless telephony.
APN	The access point name identifies the packet data network (PDN), such as the Internet, that the subscriber wants to access. When a subscriber requests access, the UE passes the requested APN to the eNodeB, which sends it to the MME for authorization. If the subscriber does not request an APN, the MME can authorize a default APN. Each PDN that the user subscribes to has an APN and an associated packet data network gateway (PGW) that the UE uses to access the PDN. The combination of APN and PGW is called a PDN subscription context. One context is the default APN, which always connects to a PDN such as the Internet unless the user activates another APN. The HSS maintains subscriber profiles, The MME uses the profile from the HSS to validate whether the subscriber is actually subscribed to the requested APN. You can also think of the APN as the set of service-level and connection parameters—such as QoS parameters—that is authorized for the UE. A given UE can access many APNs. An APN has two parts: • Network Identifier—This defines the external PDN that the user connects to through a PGW. This part of the APN is mandatory. It can be as simple as internet or have a more complicated structure such as example.net. The network identifier can optionally specify a requested subscriber service. Operator Identifier—(Optional) This defines the provider whose PDN the user connects to through a PGW. This part of the APN is often omitted. If present, it consists of the provider's Mobile Network Code (MNC) and Mobile Country Code (MCC). An APN that includes both a Network Identifier and an Operator Identifier corresponds to a DNS name for the PGW. The APN has the following format: *network-id*.mnc*mnc-number*.mcc*mcc-number.gprs* An APN can be a simple string or more complex, as shown in these examples: • fixed-wireless • web.example.net • internet.mnc99.mcc999.gprs

Table 32: Terminology for a 3GPP Fixed Wireless Access Network *(continued)*

Term	Description
Bearer	A bearer is the tunnel that connects the UE to a PDN through the PGW. A *default bearer* is established to a default PGW whenever the UE is activated. Activation means here that the UE is on and has performed authentication. A UE device has a default bearer for each PGW to which it connects. For example, if user equipment connects to the Internet through one PGW and a corporate intranet through another PGW, two default bearers will be active. Default bearers are best-effort. The UE can establish *dedicated bearers* to other PDNs that can have different QoS requirements, such as a guaranteed bit rate (GBR). Bearers encapsulate user data with GTP-U. The GTP-U information is in turn sent over a UDP connection and inside IP packets.
eNodeB	The hardware (typically in a radio tower) that connects directly to the UE over the air and to the wireless network core. Also called evolved Node B or E-UTRAN Node B. The eNodeB has the following functions: • Terminates the radio connection from the UE. • Locates the MME that authenticates the UE (SIM card) with information from the subscriber profile maintained on the HSS. • Maintains the S1-U data plane interface with the SAEGW on the BNG. An S1-U interface can support multiple eNodeBs.
GPRS	The general packet radio service is the data standard that defines the specifications that enable wireless networks to carry IP packets to external networks.
GR	The home gateway router that provides the interface between the subscriber's network and the UE. Also called a residential gateway router.
GTP	The GPRS tunneling protocol governs the creation and use of GTP tunnels that carry traffic between two GPRS support nodes (GSNs), such as an MME and an SGW. Each GTP tunnel is identified by a TEID. The receiving end of a tunnel assigns locally the TEID that the transmitting side uses. The tunnel endpoints on the nodes exchange messages to communicate the TEID values to each other.
GTP-C	The GPRS tunneling protocol, control plane. GTP-C tunnels carry packet data units and signaling messages in the control plane (S11 interface) between tunnel endpoints on the MME and the SAEGW on the BNG.
GTP-U	The GPRS tunneling protocol, user plane. GTP-U tunnels carry packet data units and signaling messages in the user (data) plane (S1-U interface) between tunnel endpoints on the eNodeB and the SAEGW on the BNG.

Table 32: Terminology for a 3GPP Fixed Wireless Access Network *(continued)*

Term	Description
HSS	The Home Subscriber Server maintains a database of subscriber and service information. This information supports call (connection) control and session management. The HSS has the following functions: • Provides authentication information from the subscriber profile to the MME. The MME uses that information to authenticate the UE for the wireless access network connection. • Identifies the APN that represents and defines the connection for the UE.
IMSI	The International Mobile Subscriber Identity number that identifies a 3GPP subscriber. The IMSI consists of a mobile country code, a mobile network code, and a mobile station identification number.
MEI	The Mobile Equipment Identity number that uniquely identifies the subscriber device.
MME	The mobility management entity is the control node for the wireless access network, communicating with eNodeB, HSS, and SAEGW. Some of its functions include the following: • Maintains the S6a interface with the HSS. • Manages and stores contexts for the UE. • Authenticates the UE with the HSS by using various types of UE identification, such as IMSI, MEI, or MSISDN. • Maintains the S11 control plane interface with the SAEGW on the BNG. An S11 interface can support multiple MMEs. • Selects the SAEGW on the BNG for the subscriber session. • Sends messages to the SAEGW for traffic control. • Participates in the bearer activation/deactivation process. • Manages the UE idle state (so the device is reachable from other devices and services), and performs idle-mode paging of UE.
MSISDN	The Mobile Subscriber ISDN Number (telephone number) that is assigned to the mobile subscriber.
PGW	The packet data network gateway provides the UE with connectivity to external networks such as the Internet. Traffic to and from the UE is processed by the PGW. The BNG functions as an SAEGW, which includes the functions of both PGW and SGW. The PGW performs the following functions: • Applies the APN characteristics to the UE session. • Allocates IP addresses to the UE during setup of the default bearer. • Filters packets to and from the subscriber. • Enforces policy. • Collects charging information for processing.

Table 32: Terminology for a 3GPP Fixed Wireless Access Network *(continued)*

Term	Description
S11	The GTPv2-based control plane interface that connects the MME and the SAEGW on the BNG. GTP-C tunnels carry control messages. An S11 interface can support multiple MMEs.
S1-MME	The GTPv2-based control plane interface that connects eNodeB and the MME.
S1-U	The GTPv1-based user plane interface that connects eNodeB and the SAEGW on the BNG. S1-U is also called the data plane interface. GTP-U tunnels on the interface carry user payloads. An S1-U interface can support multiple eNodeBs.
S6a	Interface that connects the MME and the HSS, which use this interface to exchange subscriber, service, and UE information.
SAEGW	The System Architecture Evolution gateway that includes the functions of both the SGW and the PGW. It enables the BNG to act as both SGW and PGW.
SGW	The serving gateway routes and forwards subscriber data packets. The BNG functions as an SAEGW, which includes the functions of both PGW and SGW. The SGW performs the following functions: • Terminates S1-U interfaces with eNodeBs and S11 interfaces with MMEs. • Subscriber management for UE DHCP subscribers.
TEID	A tunnel endpoint identifier that uniquely identifies a GTP tunnel endpoint in the scope of a path. A fully qualified TEID consists of an IP address concatenated with a locally allocated identifier. Four TEIDs are defined, together they uniquely identify a default bearer session: • L-TEID-C consists of the IP address of the S11 interface on the BNG and the BNG's allocated identifier. • L-TEID-U consists of the IP address of the S1-U interface on the BNG and the BNG's allocated identifier • R-TEID-C consists of the IP address of the S11 interface on the MME and the MME's allocated identifier. • R-TEID-U consists of the IP address of the S1-U interface on the eNodeB and the eNodeB's allocated identifier
UE	The user equipment that connects to the wireless network's eNodeB and to the subscriber's network. UE corresponds to what is called CPE in other contexts. In some cases, the UE consists of a SIM card and a residential gateway router (RG) that can host the SIM. In other cases the SIM might be in a separate device that connects to the RG. In both cases, the SIM provides the wireless radio connectivity to eNodeB in the fixed wireless access network.

Benefits of Fixed Wireless Access

- Reduce last-mile installation and maintenance costs by using radio backhaul towers connected to hard-wired network instead of providing fiber to the building.
- Ability to increase service offerings to underserved end users.

How to Configure Fixed Wireless Access

The fixed wireless access configuration enables the SAEGW on the BNG. At a minimum you must configure an APN, the control plane and the data plane.

The following procedure assumes that you have configured separately any of the following that apply for your APN:

- Access profile
- Dynamic profile
- Anchor point interface
- Address pool

To configure fixed wireless access on the BNG:

1. Specify the name of the pseudowire physical interface that anchors all incoming GTP-U (S1-U interface) tunnels on the BNG.

   NOTE: The BNG supports only a single anchor point.

   ```
   [edit services fixed-wireless-access]
   user@host# set anchor-point anchor-point-name
   ```

2. Define an access point name for the user equipment by specifying the connection and service parameters that the subscriber's device can use to connect to the PGW to access a PDN.

   ```
   [edit services fixed-wireless-access]
   user@host# edit apn access-point-name
   ```

   a. (Optional) Associate an authentication or accounting profile with the APN that specifies authentication or accounting parameters.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set aaa-profile aaa-profile-name
```

b. (Optional) Specify one or more authentication parameters that the BNG sends to the external AAA server for subscribers using this APN. You configure details about the external server in the access profile that you associate with the APN. Some networks might not use this authentication, because the HSS has already authenticated the UE and determined subscriber access. Other networks might not use this because they use Diameter for online charging.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set authentication password password
user@host# set authentication username-include delimiter delimiter
user@host# set authentication username-include domain-name domain-name
user@host# set authentication username-include imsi
user@host# set authentication username-include mei
user@host# set authentication username-include msisdn
user@host# set authentication username-include user-prefix user-prefix
```

c. (Optional) Specify the type of data sent to the APN.

NOTE: Only the **ipv4** type is supported.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set apn-data-type type
```

d. (Optional) Specify a text description for the APN. You can use this to provide more information about the APN than its name alone can convey. The description for an APN appears in subscriber profiles in the HSS database.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set description description
```

e. Associate a dynamic profile with the APN to create the dynamic fixed wireless interface.

> **NOTE:** Services such as CoS and firewall filters are applied as part of the DHCP configuration.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set dynamic-profile dynamic-profile
```

f. (Optional) Specify the name of the address pool for IPv4 addresses assigned to the APN.

> **NOTE:** Only IPv4 addresses are supported.

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set ipv4-address-pool pool-name
```

g. (Optional)

```
[edit services fixed-wireless-access apn access-point-name]
user@host# set routing-instance routing-instance-name
```

3. Configure the control plane by specifying the name of the MME and the IPv4 address of the S11 interface on the BNG. The S11 interface is the reference point or connection between the MME and the SAEGW on the BNG for control packets.

```
[edit services fixed-wireless-access]
user@host# set control-plane name s11 v4-address ip-address
```

4. Configure the data plane by specifying the name of the eNodeB and the IPv4 address for the S1-U interface on the BNG. The S1-U interface is the reference point or connection between the eNodeB and the SAEGW on the BNG for subscriber data traffic.

```
[edit services fixed-wireless-access]
user@host# set data-plane name s1-u v4-address ip-address
```

For example, the following configuration specifies the S11 interface address as 192.0.2.30 and the S1-U interface address as 192.0.2.100. The pseudowire anchor is ps0. The APN is named internet-basic; it has

both an access profile for RADIUS parameters and a dynamic client profile attached. The authentication password is $ABC123. The username includes the domain name example.net and the subscriber's IMSI. The data type is the required IPv4. A descriptive string, fwa-basic-subscribers-phoenix, is associated with the APN.

The APN uses the default routing instance, because no other routing instance is configured. IP addresses are supplied by RADIUS, because the configuration does not specify a local pool.

```
fixed-wireless-access {
    anchor-point ps0;
    control-plane mme-45 {
        s11 {
            v4-address ip-address;
        }
    }
    data-plane x-s1u-20 {
        s1-u {
            v4-address 192.0.2.100;
        }
    }
    apn internet-basic {
        aaa-profile fwa-radius-prof;
        apn-data-type ipv4
        authentication {
            password $ABC123;
            username-include {
                domain-name example.net;
                imsi;
            }
        }
        description fwa-basic-subscribers-phoenix;
        dynamic-profile fwa-dyn-profile;
    }
}
```

The dynamic profile that you attach to the APN creates the dynamic interface for fixed wireless access. The following configuration is a simple example:

```
dynamic-profiles fwa-dyn-profile {
    interfaces {
        "$junos-interface-ifd-name" {
            unit "$junos-interface-unit" {
                family inet {
                    unnumbered-address lo0.0 preferred-source-address 10.0.0.1;
```

```
            }
          }
        }
      }
    }
```

The anchor point interface configuration is also very simple. For example:

```
interfaces {
    ps0 {
        anchor-point {
            rlt-0/1/10;
        }
        flexible-vlan-tagging;
    }
}
```

Verifying and Monitoring Fixed Wireless Access

Purpose

Determine status information and statistics for fixed wireless access configurations.

Action

- To display a list of all interfaces on the BNG supporting fixed wireless access subscribers:

  ```
  user@host>show subscribers client-type fixed-wireless-access
  ```

- To display detailed information about fixed wireless access subscribers, including username and IP address, dynamic profile, local and remote TEIDs, and remote and local IP addresses for the BNG connection to eNodeBs and MMEs:

  ```
  user@host>show subscribers client-type fixed-wireless-access detail
  ```

- To view statistics for messages exchanged between the BNG, eNodeB, and MME:

  ```
  user@host>show services fixed-wireless-access statistics
  ```

Release History Table

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, MX Series routers acting as BNGs can support subscribers in Third-Generation Partnership Project (3GPP) fixed wireless access networks, enabling the integration of fixed wireless subscribers with the wireline subscriber management backend.

RELATED DOCUMENTATION

Tracing Fixed Wireless Access Events for Troubleshooting | 388

Configuring Access Profile Options for Interactions with RADIUS Servers

Dynamic Profiles for Subscriber Management

Address-Assignment Pools for Subscriber Management

Tracing Fixed Wireless Access Events for Troubleshooting

IN THIS SECTION

- Configuring the Fixed Wireless Access Trace Log Filename | 389
- Configuring the Number and Size of Fixed Wireless Access Log Files | 389
- Configuring Access to the Fixed Wireless Access Log File | 390
- Configuring a Regular Expression for Fixed Wireless Access Messages to Be Logged | 391
- Configuring the Fixed Wireless Access Tracing Flags | 391

The Junos OS trace feature tracks fixed wireless access 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 **tcpfwdd**. 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.

The following topics describe how to configure all aspects of tracing fixed wireless access operations:

Configuring the Fixed Wireless Access Trace Log Filename

By default, the name of the file that records trace output for fixed wireless access is **bbe-fwsd**. You can specify a different name with the **file** option.

To configure the filename for fixed wireless access tracing operations:

- Specify the name of the file used for the trace output.

```
[edit system processes tcp-forwarding traceoptions]
user@host# set file fwsd_1
```

Configuring the Number and Size of Fixed Wireless Access 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 system processes tcp-forwarding traceoptions]
user@host# set file fwsd_1 _logfile_1 files 20 size 2097152
```

Configuring Access to the Fixed Wireless Access 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 system processes tcp-forwarding traceoptions]
user@host# set file fwsd_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 system processes tcp-forwarding traceoptions]
user@host# set file fwsd_1 _logfile_1 no-world-readable
```

Configuring a Regular Expression for Fixed Wireless Access Messages to Be Logged

By default, the trace operation output includes all messages 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 system processes tcp-forwarding traceoptions]
user@host# set file fwsd_1 _logfile_1 match regex
```

Configuring the Fixed Wireless Access 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 system processes tcp-forwarding traceoptions]
user@host# set flag flag-name
```

RELATED DOCUMENTATION

Fixed Wireless Access Networks | 374

8

CHAPTER

Configuration Statements

aaa-access-profile (L2TP LNS) | **401**

aaa-context (AAA Options) | **402**

aaa-options (Access Profile) | **403**

aaa-options (PPP Profile) | **404**

access (Dynamic Access Routes) | **406**

access-internal (Dynamic Access-Internal Routes) | **408**

access-line (Access-Line Rate Adjustment) | **410**

access-line-information (L2TP) | **426**

access-profile (AAA Options) | **428**

address (L2TP Destination) | **429**

address (L2TP Tunnel Destination) | **430**

address (LNS Local Gateway) | **431**

address (Tunnel Profile Remote Gateway) | **432**

address (Tunnel Profile Source Gateway) | **433**

address-change-immediate-update | **434**

aggregated-inline-services-options (Aggregated Inline Services) | **435**

allow-snooped-clients | **437**

always-write-option-82 | **439**

anchor-point (Pseudowire Subscriber Interfaces) | **441**

assignment-id-format (L2TP LAC) | **444**

authentication (Static and Dynamic PPP) | **445**

avp (L2TP Tunnel Switching) | **446**

bandwidth (Inline Services) | **447**

bandwidth (Tunnel Services) | **449**

bearer-type (L2TP Tunnel Switching) | **451**

bfd | **452**

calling-number (L2TP Tunnel Switching) | **454**

challenge-length (Static and Dynamic PPP) | **455**

chap | **457**

chap (Dynamic PPP) | **459**

chap (L2TP) | **460**

cisco-nas-port-info (L2TP Tunnel Switching) | **461**

client | **462**

delimiter (Access Profile) | **466**

destination (L2TP) | **467**

destination-equal-load-balancing (L2TP LAC) | **468**

destruct-timeout (L2TP) | **469**

detection-time | **470**

device-count (Pseudowire Subscriber Interfaces) | **471**

dhcp-local-server | **472**

dhcp-relay | **489**

dhcpv6 (DHCP Local Server) | **508**

dhcpv6 (DHCP Relay Agent) | **517**

dial-options | **526**

dial-options (Dynamic Profiles) | **528**

disable-calling-number-avp (L2TP LAC) | **529**

disable-failover-protocol (L2TP) | **530**

drain | **532**

dual-stack-group (DHCP Local Server) | **533**

dual-stack-group (DHCP Relay Agent) | **535**

duplicate-clients (DHCPv6 Local Server and Relay Agent) | **538**

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent) | **540**

dynamic-profile (L2TP) | **542**

dynamic-profile (PPP) | **543**

dynamic-profiles | **544**

enable-ipv6-services-for-lac (L2TP) | **558**

enable-snmp-tunnel-statistics (L2TP) | **559**

enforce-strict-scale-limit-license (Subscriber Management) | **560**

equals (Dynamic Profile) | **561**

failover-resync | **562**

failover-within-preference (L2TP LAC) | **563**

failure-action | **564**

flexible-vlan-tagging | **566**

forward-snooped-clients (DHCP Local Server) | **567**

forward-snooped-clients (DHCP Relay Agent) | **568**

fpc (MX Series 5G Universal Routing Platforms) | **570**

gateway-name (LNS Local Gateway) | **572**

gateway-name (Tunnel Profile Remote Gateway) | **573**

gateway-name (Tunnel Profile Source Gateway) | **574**

gres-route-flush-delay (Subscriber Management) | **575**

group (DHCP Local Server) | **576**

group (DHCP Relay Agent) | **582**

group-profile (Group Profile) | **588**

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | **590**

holddown-interval | **592**

hello-interval (L2TP) | **593**

identification (Tunnel Profile) | **594**

idle-timeout (Access) | **595**

idle-timeout (L2TP) | **597**

ignore-magic-number-mismatch (Access Group Profile) | **598**

ignore-magic-number-mismatch (Dynamic Profiles) | **600**

initiate-ncp (Dynamic and Static PPP) | **602**

inline-services (PIC level) | **604**

input-hierarchical-policer | **605**

interface (Dynamic Routing Instances) | **606**

interface (Service Interfaces) | **607**

interface-id | **608**

interfaces (Static and Dynamic Subscribers) | **609**

ip-reassembly | **616**

ip-reassembly (L2TP) | **617**

ip-reassembly-rules (Service Set) | **618**

ipcp-suggest-dns-option | **619**

keepalive | **620**

keepalives | **621**

keepalives (Dynamic Profiles) | **623**

l2tp | **625**

l2tp (Profile) | **629**

l2tp-access-profile | **633**

l2tp-maximum-session (Service Interfaces) | **634**

layer2-liveness-detection (Receive) | **636**

layer2-liveness-detection (Send) | **638**

lcp-renegotiation | **640**

liveness-detection | **641**

local-authentication (Dynamic PPP Options) | **643**

local-gateway (L2TP LNS) | **644**

lockout-timeout (L2TP Destination Lockout) | **645**

logical-system (Tunnel Profile) | **646**

mac | **647**

mac-address (Dynamic Access-Internal Routes) | **648**

match-direction (IP Reassembly Rule) | **649**

maximum-sessions (L2TP) | **650**

maximum-sessions-per-tunnel | **651**

max-sessions (Tunnel Profile) | **652**

medium (Tunnel Profile) | **653**

method | **654**

metric (Dynamic Access-Internal Routes) | **656**

minimum-interval | **657**

minimum-receive-interval | **659**

minimum-retransmission-timeout (L2TP Tunnel) | **660**

mtu | **661**

multiplier | **665**

name (L2TP Destination) | **666**

name (L2TP Tunnel Destination) | **667**

no-adaptation | **668**

nas-port-method (L2TP LAC) | **669**

nas-port-method (Tunnel Profile) | **670**

next-hop (Dynamic Access Routes) | **671**

next-hop-service | **673**

no-allow-snooped-clients | **675**

no-gratuitous-arp-request | **676**

no-snoop (DHCP Local Server and Relay Agent) | **677**

on-demand-ip-address | **679**

options (Access Profile) | **681**

override (RADIUS Options) | **692**

overrides (DHCP Relay Agent) | **694**

overrides (Enhanced Subscriber Management) | **696**

pap | **700**

pap (Dynamic PPP) | **701**

pap (L2TP) | **702**

parse-direction (Access Profile) | **703**

pic (M Series and T Series Routers) | **705**

pool (Service Interfaces) | **707**

pp0 (Dynamic PPPoE) | **708**

ppp (Group Profile) | **710**

ppp-options | **713**

ppp-options (Dynamic PPP) | **715**

ppp-options (L2TP) | **718**

preference (Subscriber Management) | **720**

preference (Tunnel Profile) | **721**

primary-interface (Aggregated Inline Services) | **722**

profile (Access) | **723**

proxy-mode | **730**

ps0 (Pseudowire Subscriber Interfaces) | **731**

pseudowire-service (Pseudowire Subscriber Interfaces) | **732**

qualified-next-hop (Dynamic Access-Internal Routes) | **733**

radius (Access Profile) | **734**

reject-unauthorized-ipv6cp | **738**

relay-option-82 | **739**

remote-gateway (Tunnel Profile) | **741**

report-ingress-shaping-rate (Dynamic CoS Interfaces) | **742**

request services l2tp destination unlock | **744**

retransmission-count-established (L2TP) | **745**

retransmission-count-not-established (L2TP) | **746**

route (Access) | **747**

route (Access Internal) | **749**

route-suppression (DHCP Local Server and Relay Agent) | **751**

routing-instance (Tunnel Profile) | **752**

routing-instance (L2TP Destination) | **753**

routing-instance (L2TP Tunnel Destination) | **754**

routing-instances (Dynamic Profiles) | **755**

routing-options (Dynamic Profiles) | **757**

rule (IP Reassembly) | **759**

rx-connect-speed-when-equal (L2TP LAC) | **760**

rx-window-size (L2TP) | **761**

secondary-interface (Aggregated Inline Services) | **762**

secret (Tunnel Profile) | **763**

service-device-pool (L2TP) | **764**

service-device-pools (Service Interfaces) | **765**

service-interface (L2TP Processing) | **766**

service-profile (L2TP) | **767**

service-rate-limiter (Access) | **769**

session-mode | **771**

session-options | **772**

sessions-limit-group (L2TP) | **776**

soft-gre | **777**

source-gateway (Tunnel Profile) | **779**

stacked-vlan-tagging | **780**

statistics (Access Profile) | **781**

strip-user-name (Access Profile) | **782**

subscriber-context (AAA Options) | **783**

subscriber-management (Subscriber Management) | **784**

tag (Access) | **787**

tag2 (Dynamic Access Routes) | **788**

threshold (detection-time) | **789**

threshold (transmit-interval) | **791**

tos-reflect (L2TP) | **792**

trace (DHCP Relay Agent) | **793**

traceoptions (Services L2TP) | **795**

traceoptions (Protocols PPP Service) | **800**

traceoptions (Subscriber Management) | **804**

transmit-interval | **806**

tunnel (L2TP) | **807**

tunnel (Tunnel Profile) | **809**

tunnel-group | **811**

tunnel-profile (L2TP Tunnel Switching) | **813**

tunnel-profile (Tunnel Profile) | **814**

tunnel-switch-profile (L2TP Tunnel Switching, Application) | **815**

tunnel-switch-profile (L2TP Tunnel Switching, Definition) | **816**

tx-address-change (L2TP LAC) | **817**

tx-connect-speed-method (L2TP LAC) | **819**

type (Tunnel Profile) | **822**

unit (Dynamic PPPoE) | **823**

unit (Dynamic Profiles Standard Interface) | **826**

untagged | **830**

username-include (Local Authentication) | **831**

version (BFD) | **833**

weighted-load-balancing (L2TP LAC) | **834**

vlan-id (Dynamic Profiles) | **835**

vlan-tagging | **836**

vlan-tagging (Dynamic) | **839**

vlan-tags | **840**

aaa-access-profile (L2TP LNS)

Syntax

```
aaa-access-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name],
[edit access profile profile-name client client-name l2tp]
```

Release Information

Statement introduced in Junos OS Release 11.4.
Support at the **[edit access profile *profile-name* client *client-name* l2tp]** 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

profile-name—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

- Configuring an L2TP LNS with Inline Service Interfaces | **261**
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**
- Configuring an L2TP Access Profile on the LNS | **268**

aaa-context (AAA Options)

Syntax

```
aaa-context aaa-context-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

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.

NOTE: Only the default (master) logical system is supported.

Options

aaa-context-name—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

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

aaa-options (Access Profile)

Syntax

```
aaa-options aaa-options-name {
    aaa-context aaa-context-name;
    access-profile profile-name;
    subscriber-context subscriber-context-name
}
```

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

aaa-options-name—Name of the set of options.

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

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Username Modification for Subscriber Sessions*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

aaa-options (PPP Profile)

Syntax

```
aaa-options aaa-options-name;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp ppp-options],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

Specify a set of AAA options that is used for authentication of PPP subscribers. The set of options is defined globally with the **aaa-options** ***aaa-options-name*** statement at the **[edit access]** hierarchy level.

You can specify the option set in a dynamic PPP profile or in a group profile.

- 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 **access-profile** ***profile-name*** 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.

> **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.
>
> This meant that **aaa-options** configured in a group profile is not applied when the dynamic profile includes any PPP-option, even when the dynamic profile does not include **aaa-options**.

Options

aaa-options-name—Name of the set of options.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Username Modification for Subscriber Sessions*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

access (Dynamic Access Routes)

Syntax

```
access {
  route prefix {
    next-hop next-hop;
    metric route-cost;
    preference route-distance;
    tag route-tag;
    tag2 route-tag2;
  }
```

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 in a dynamic client profile.

> **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. 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.

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 | 64

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 in a dynamic client profile. 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. 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.

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 and PPP Subscribers | 66

access-line (Access-Line Rate Adjustment)

List of Syntax

Syntax for Releases Earlier than Junos OS Release 19.3R1 on page 410
Syntax for Junos OS Release 19.3R1 and Higher Releases. on page 412

Syntax for Releases Earlier than Junos OS Release 19.3R1

```
access-line {
    adsl-overhead-bytes bytes;
    adsl-total-adjust percentage;
    adsl2-overhead-bytes bytes;
    adsl2-total-adjust percentage;
    adsl2-plus-overhead-bytes bytes;
    adsl2-plus-total-adjust percentage;
    gfast-bonded-overhead-adjust percentage
    gfast-bonded-overhead-bytes bytes
    gfast-bonded-total-adjust percentage
    gfast-overhead-adjust percentage
    gfast-overhead-bytes bytes
    gfast-total-adjust percentage
    hierarchical-access-network-detection;
    other-overhead-adjust percentage;
    other-overhead-bytes bytes;
    other-total-adjust percentage;
    sdsl-bonded-overhead-bytes  bytes
    sdsl-bonded-overhead-adjust  percentage
    sdsl-bonded-total-adjust  percentage
    sdsl-overhead-adjust percentage;
    sdsl-overhead-bytes bytes;
    sdsl-total-adjust percentage;
    vdsl-overhead-adjust percentage;
    vdsl-overhead-bytes bytes;
    vdsl-total-adjust percentage;
    vdsl2-annex-q-bonded-overhead-adjust  percentage
    vdsl2-annex-q-bonded-overhead-bytes bytes
    vdsl2-annex-q-bonded-total-adjust  percentage
    vdsl2-annex-q-overhead-adjust  percentage
    vdsl2-annex-q-overhead-bytes  bytes
    vdsl2-annex-q-total-adjust  percentage
    vdsl2-bonded-overhead-adjust  percentage
    vdsl2-bonded-overhead-bytes  bytes
    vdsl2-bonded-total-adjust  percentage
    vdsl2-overhead-adjust percentage;
```

```
    vdsl2-overhead-bytes bytes;
    vdsl2-total-adjust percentage;
}
```

Syntax for Junos OS Release 19.3R1 and Higher Releases.

```
access-line {
    attributes {
        preference (dsl | pon);
    }
    dsl {
        adsl {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2 {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2-plus {
            overhead-bytes bytes;
            total-adjust percent;
        }
        gfast {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        gfast-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        type tlv-value {
```

```
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    vdsl {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    vdsl2 {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    vdsl2-annex-q {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    vdsl2-annex-q-bonded {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
}
hierarchical-access-network-detection;
```

```
    pon {
        gpon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        twdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        type tlv-value {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        wdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xg-pon1 {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xgs-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
    }
}
```

Hierarchy Level

```
[edit system]
```

Release Information

Statement introduced in Junos OS Release 17.4R1.

The following options added in Junos OS 18.2R1:

gfast-bonded-overhead-adjust	**vdsl2-annex-q-bonded-overhead-adjust**
gfast-bonded-overhead-bytes	**vdsl2-annex-q-bonded-overhead-bytes**
gfast-bonded-total-adjust	**vdsl2-annex-q-bonded-total-adjust**
gfast-overhead-adjust	**vdsl2-annex-q-overhead-adjust**
gfast-overhead-bytes	**vdsl2-annex-q-overhead-bytes**
gfast-total-adjust	**vdsl2-annex-q-total-adjust**
sdsl-bonded-overhead-bytes	**vdsl2-bonded-overhead-bytes**
sdsl-bonded-overhead-adjust	**vdsl2-bonded-total-adjust**
sdsl-bonded-total-adjust	**vdsl2-bonded-total-adjust**

hierarchical-access-network-detection option added in Junos OS 18.4R1.

attributes, **dsl**, and **pon** statements added in Junos OS Release 19.3R1.

Description

Configure values to adjust data rates by a percentage of the actual data rate, or adjust encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes. Depending on the value, it may be reported to AAA, CoS, or both.

The actual (unadjusted) downstream and upstream data rates, line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

Adjustments are applied to all subscribers using access lines of the specific subscriber access line type:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following messages and AVPs:

- AVP 24, Tx Connect Speed (ICCN message)
- AVP 38, Rx Connect Speed (ICCN message)
- AVP 97, Connect Speed Update (CSUN message)

NOTE: Starting in Junos OS Release 19.3R1, the pre-existing adjustment options were renamed and placed in the new **dsl** stanza. The old DSL options are deprecated, but they redirect to the new location. The **hierarchical-access-network-detection** option is unchanged.

BEST PRACTICE: We recommend that you update your scripts to use the *dsl* statement when you upgrade to Junos OS Release 19.3R1 or higher releases. The redirect function will be supported for only a limited time.

Options

adsl-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

> **NOTE:** This option replaces the **adsl-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

adsl-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL access line. The adjusted rate is reported only to AAA.

> **NOTE:** This option replaces the **qos-adjust-adsl** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

adsl2-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

> **NOTE:** This option replaces the **adsl2-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

adsl2-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2 access line. The adjusted rate is reported only to AAA.

NOTE: This option replaces the **qos-adjust-adsl2** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

adsl2-plus-overhead-bytes ***bytes***—Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2+ access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

NOTE: This option replaces the **adsl2-plus-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

adsl2-plus-total-adjust ***percentage***—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2+ access line. The adjusted rate is reported only to AAA.

NOTE: This option replaces the **qos-adjust-adsl2-plus** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

gfast-bonded-overhead-adjust ***percentage***—Adjustment factor in percent that is applied to the downstream and upstream bonded data overhead rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

Range: 80 through 100 percent

Default: 100 percent

gfast-bonded-overhead-bytes ***bytes***—Number of bytes added to or subtracted from the actual downstream cell bonded overhead for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure. Specify G.fast bonded value in bytes.

Range: -100 through 100 bytes

Default: 0 bytes

gfast-bonded-total-adjust *percentage*—Adjustment factor in percent that is globally applied to the downstream and upstream bonded data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

Range: 0 through 100 percent

Default: 100 percent

gfast-overhead-adjust *percentage*—Adjustment factor in percent that is applied to the downstream and upstream data overhead rates for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

Range: 80 through 100 percent

Default: 100 percent

gfast-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

Range: -100 through 100 bytes

Default: 0 bytes

gfast-total-adjust *percentage*—Adjustment factor in percent that is globally applied to the downstream and upstream data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

Range: 0 through 100 percent

Default: 100 percent

hierarchical-access-network-detection—Enable parsing of ANCP subscriber access loop attributes (TLVs) for backhaul line identifiers, to recognize when the access node references a logical interface set rather than an individual subscriber. If the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x03) string begins with a **#** sign, then the remainder of the string represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. The string is used as the name of a CoS Level 2 interface set that groups subscribers.

NOTE: The Access-Loop-Remote-ID (TLV (0x02) is similarly parsed for the **#** character, but the resulting string is not used in the current release.

These TLVs can be parsed in ANCP messages or PPPoE IA tags in PADR messages.

Default: Disabled, in case some users include an initial **#** character for some other purpose.

other-overhead-adjust *percentage*—Adjust the actual downstream rate for an access line of DSL type OTHER by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

NOTE: This option replaces the **other-overhead-adjust** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: 80 through 100 percent

Default: 100 percent

other-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an access line of DSL type OTHER to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

NOTE: This option replaces the **other-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

other-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an access line of DSL type OTHER. The adjusted rate is reported only to AAA.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

NOTE: This option replaces the **qos-adjust-other** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

sdsl-bonded-overhead-adjust *percentage*—Adjust the actual downstream rate for an SDSL bonded access line by multiplying the rate by the specified percentage.

Range: 80 through 100 percent

Default: 100 percent

sdsl-bonded-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL bonded access line to account for the traffic encapsulation overhead.

Range: -100 through 100 bytes

Default: 0 bytes

sdsl-bonded-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL bonded access line. This value is reported to AAA.

Range: 0 through 100 percent

Default: 100 percent

sdsl-overhead-adjust *percentage*—Adjust the actual downstream rate for an SDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

NOTE: This option replaces the **sdsl-overhead-adjust** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: 80 through 100 percent

Default: 100 percent

sdsl-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

NOTE: This option replaces the **sdsl-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

sdsl-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL access line. The adjusted rate is reported only to AAA.

> **NOTE:** This option replaces the **qos-adjust-sdsl** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

vdsl-overhead-adjust *percentage*—Adjust the actual downstream rate for a VDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

> **NOTE:** This option replaces the **vdsl-overhead-adjust** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: 80 through 100 percent

Default: 100 percent

vdsl-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

> **NOTE:** This option replaces the **vdsl-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

vdsl-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL access line. The adjusted rate is reported only to AAA.

> **NOTE:** This option replaces the **qos-adjust-vdsl** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

vdsl2-annex-q-bonded-overhead-adjust *percentage*—Adjust the actual downstream rate for a VDSL2 annex q bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

Range: 80 through 100 percent

Default: 100 percent

vdsl2-annex-q-bonded-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q bonded access line to account for the traffic encapsulation overhead.

Range: -100 through 100 bytes

Default: 0 bytes

vdsl2-annex-q-bonded-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q bonded access line. The adjusted rate is reported to AAA.

Range: 0 through 100 percent

Default: 100 percent

vdsl2-annex-q-overhead-adjust *percentage*—Adjust the actual downstream rate for a VDSL2 annex q access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

Range: 80 through 100 percent

Default: 100 percent

vdsl2-annex-q-overhead-bytes *bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q access line to account for the traffic encapsulation overhead.

Range: -100 through 100 bytes

Default: 0 bytes

vdsl2-annex-q-total-adjust ***percentage***—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q access line. The adjusted rate is reported to AAA.

vdsl2-bonded-overhead-adjust ***percentage***—Adjust the actual downstream rate for a VDSL2 bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

Range: 80 through 100 percent

Default: 100 percent

vdsl2-bonded-overhead-bytes*bytes*—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 bonded access line to account for the traffic encapsulation overhead.

Range: -100 through 100 bytes

Default: 0 bytes

vdsl2-bonded-total-adjust ***percentage***—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 bonded access line. The adjusted rate is reported to AAA.

Range: 0 through 100 percent

Default: 100 percent

vdsl2-overhead-adjust ***percentage***—Adjust the actual downstream rate for a VDSL2 access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

NOTE: This option replaces the **vdsl2-overhead-adjust** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: 80 through 100 percent

Default: 100 percent

vdsl2-overhead-bytes ***bytes***—Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the **qos-adjust** statement at the **[edit protocols ancp]** hierarchy level.

Number of bytes added to or subtracted from the actual downstream frame overhead.

NOTE: This option replaces the **vdsl2-bytes** statement at the **[edit protocols ancp qos-adjust]** hierarchy level.

Range: -100 through 100 bytes

Default: 0 bytes

vdsl2-total-adjust *percentage*—Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 access line. The adjusted rate is reported only to AAA.

> **NOTE:** This option replaces the **qos-adjust-vdsl2** statement at the **[edit protocols ancp]** hierarchy level.

Range: 0 through 100 percent

Default: 100 percent

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 the ANCP Agent to Report Traffic Rates to CoS*
- *Traffic Rate Reporting and Adjustment by the ANCP Agent*
- *Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates*
- *Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates*
- *Configuring the ANCP Agent*

access-line-information (L2TP)

Syntax

```
access-line-information <connection-speed-update>;
```

Hierarchy Level

```
[edit services l2tp],
[edit services l2tp destination ip-address]
```

Release Information

Statement introduced in Junos OS Release 14.1.
Support at the **[edit services l2tp]** 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

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.

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.

Including this statement at the **[edit services l2tp destination *ip-address*]** 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.

This statement has no effect on existing subscribers; it applies only to new subscribers.

Options

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.

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.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Reporting and Processing of Subscriber Access Line Information | 247

access-profile (AAA Options)

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

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

profile-name—Name of the subscriber access profile that includes a subscriber username stripping configuration.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Subscriber Session Timeout Options

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

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. 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

Configuring L2TP Drain | 181

address (L2TP Tunnel Destination)

Syntax

```
address ip-address; {
    drain;
    routing-instance routing-instance-name {
        drain;
    }
}
```

Hierarchy Level

```
[edit services l2tp tunnel name name ]
```

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

ip-address—IP address of the tunnel destination.

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

Configuring L2TP Drain | 181

address (LNS Local Gateway)

Syntax

address *address*;

Hierarchy Level

[edit services l2tp tunnel-group *name* **local-gateway**]

Release Information
Statement introduced before Junos OS Release 7.4.

Description

Specify the local (LNS) IP address for L2TP tunnel.

Options

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.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Local Gateway Address and PIC.

Configuring L2TP Tunnel Groups

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

address (Tunnel Profile Remote Gateway)

Syntax

address *server-ip-address*;

Hierarchy Level

[edit access tunnel-profile *profile-name* tunnel *tunnel-id* **remote-gateway**]

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

server-ip-address—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

Configuring a Tunnel Profile for Subscriber Access | 216

address (Tunnel Profile Source Gateway)

Syntax

```
address client-ip-address;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id source-gateway]
```

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

client-ip-address—IP address of the source 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

Configuring a Tunnel Profile for Subscriber Access | 216

address-change-immediate-update

Syntax

```
address-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Release Information
Statement introduced in Junos OS Release 13.1.

Description
Configure the router to send an Interim-Accounting message to the RADIUS server immediately after on-demand IPv4 allocation and de-allocation.

Changes to this setting take effect for new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Default
This functionality is disabled by default.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *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. 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

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278

Configuring an L2TP LNS with Inline Service Interfaces | 261

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 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*
- *Overriding the Default DHCP Relay Configuration Settings*
- *DHCP Snooping Support*
- *Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent*

always-write-option-82

Syntax

```
always-write-option-82;
```

Hierarchy Level

[edit forwarding-options dhcp-relay **overrides**],
[edit forwarding-options dhcp-relay group *group-name* **overrides**],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay **overrides**],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay **overrides**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* **overrides**]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

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:

- 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

Using DHCP Relay Agent Option 82 Information

Extended DHCP Relay Agent Overview

anchor-point (Pseudowire Subscriber Interfaces)

Syntax

```
anchor-point lt-device;
```

Hierarchy Level

```
[edit interfaces ps0]
```

Release Information
Statement introduced in Junos OS Release 13.1.

Description
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 that hosts subscriber management and enables you to perform subscriber management services at the 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).

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
   ```

NOTE: You cannot disable the underlying logical tunnel (lt) interface or redundant logical tunnel (rlt) interface when a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

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. Disable the underlying interface and commit.

```
[edit interfaces]
user@host# set interfaces underlying-interface-name disable
user@host# commit
```

Options

lt-device—An lt device in the format **lt-*fpc/pic/port***

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 | 331
- Configuring a Pseudowire Subscriber Logical Interface | 338
- Configuring a Pseudowire Subscriber Logical Interface Device | 341
- Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview | 334

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.

> **NOTE:** Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp tunnel assignment-id-format**.

Default
assignment-Id

Options
assignment-Id—The tunnel name corresponds to RADIUS attribute Tunnel-Assignment-Id [82].

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].

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 | 215

authentication (Static and Dynamic PPP)

Syntax

```
authentication [ authentication-protocols ];
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit unit-number ppp-options]
```

Release Information
Statement introduced in Junos OS Release 12.2.

Description
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.

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 ([ ]).

Options
authentication-protocols—One or both of the following PPP authentication protocols:

- **chap**—Challenge Handshake Authentication Protocol
- **pap**—Password Authentication Protocol

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Controlling the Negotiation Order of PPP Authentication Protocols | 140

avp (L2TP Tunnel Switching)

Syntax

```
avp {
    bearer-type;
    calling-number;
    cisco-nas-port-info;
}
```

Hierarchy Level

[edit access **tunnel-switch-profile** *profile-name*]

Release Information
Statement introduced in Junos OS Release 13.2.

Description
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.

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
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring L2TP Tunnel Switching | 175

bandwidth (Inline Services)

Syntax

bandwidth *bandwidth-value*;

Hierarchy Level

[edit chassis **fpc** *slot-number* pic *number* **inline-services**]

Release Information

Statement introduced in Junos OS Release 11.4.
20g, **30g**, and **40g** options added in Junos OS Release 14.1R3.
100g option added in Junos OS Release 18.2R1 on MX Series Routers with MPC7E, MPC8E, and MPC9E line cards.
50g, **60g**, **70g**, **80g**, **90g**, **200g**, **300g** and **400g** options added in Junos OS Release 18.3R1.

Description

Configure the amount of bandwidth in gigabits per second reserved on each Packet Forwarding Engine for tunnel traffic using inline services. Configuring the bandwidth creates a virtual tunnel interface that is represented as **si-<*fpc*/*pic*/*port*>**.

Starting in Junos OS Release 16.2, you are not longer required to explicitly specify a bandwidth for L2TP LNS tunnel traffic using inline services. When you do not specify a bandwidth, the maximum bandwidth supported on the PIC is automatically available for the inline services; inline services can use up to this maximum value. In earlier releases, you must specify a bandwidth when you enable inline services with the **inline-services** statement.

Options

bandwidth-value—Amount of bandwidth in Gbps to reserve for tunnel traffic using inline services. You can configure bandwidth values can be as follows:

- **1g**
- **10g** through **100g** in 10 Gbps increments: **10g**, **20g**, **30g**, **40g**, **50g**, **60g**, **70g**, **80g**, **90g**, **100g**
- **100g** through **400g** in 100 Gbps increments: **100g**, **200g**, **300g**, **400g**

NOTE: Values of 100 Gbps and up are available only on MPC7E, MPC8E, and MPC9E line cards.

Default: Maximum bandwidth supported on the PIC.

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 | 273

Configuring an L2TP LNS with Inline Service Interfaces | 261

bandwidth (Tunnel Services)

Syntax

```
bandwidth bandwidth-value;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number tunnel-services]
```

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 5G Universal Routing Platforms and T4000 Core Routers only) Configure the amount of bandwidth in gigabits per second reserved on each Packet Forwarding Engine for tunnel traffic using tunnel services. Configuring the bandwidth creates a virtual tunnel interface that is represented as **lt-*<fpc/pic/port>***.

Options

bandwidth-value—Amount of bandwidth in Gbps to reserve for tunnel traffic using tunnel services:

- On ACX Series routers, the bandwidth values can be **1g** or **10g**.
- On MX Series routers, the bandwidth values can be as follows:
  - **1g**
  - **10g** through **100g** in 10 Gbps increments: **10g**, **20g**, **30g**, **40g**, **50g**, **60g**, **70g**, **80g**, **90g**, **100g**
  - **100g** through **400g** in 100 Gbps increments: **100g**, **200g**, **300g**, **400g**
- On T4000 routers, the bandwidth values can be **10g** through **100g** in 10 Gbps increments: **10g**, **20g**, **30g**, **40g**, **50g**, **60g**, **70g**, **80g**, **90g**, **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

- *Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC*
- *Tunnel Interface Configuration on MX Series Routers Overview*
- *Configuring Tunnel Interfaces on T4000 Routers*
- *Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC*
- *Example: Configuring Tunnel Interfaces on the MPC3E*
- *tunnel-services (Chassis)*

bearer-type (L2TP Tunnel Switching)

Syntax

```
bearer-type action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name 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

action—One of the following actions:

- **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.

Default: relay

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 | 175

bfd

Syntax

```
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;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit system services dhcp-local-server group group-name liveness-detection method],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method],
[edit forwarding-options dhcp-relay group group-name liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method]
```

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) as the liveness detection method.

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

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

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

calling-number (L2TP Tunnel Switching)

Syntax

```
calling-number action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name 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

action—One of the following actions:

- **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.

Default: relay

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 | 175

challenge-length (Static and Dynamic PPP)

Syntax

```
challenge-length minimum minimum-length maximum maximum-length;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options
  chap],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options chap],
[edit interfaces pp0 unit unit-number 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.

> **BEST PRACTICE:** We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.

Options

minimum-length—Minimum length, in bytes, of the CHAP challenge.

Range: 8 through 63

Default: 16

maximum-length—Maximum length, in bytes, of the CHAP challenge. The *maximum-length* must be equal to or greater than the *minimum-length*.

Range: 8 through 63

Default: 32

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

chap

Syntax

```
chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* **ppp-options**],
[edit interfaces *interface-name* unit *logical-unit-number* **ppp-options**],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **ppp-options**]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

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.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.

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** statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.

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

Configuring the PPP Challenge Handshake Authentication Protocol

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

chap (Dynamic PPP)

Syntax

```
chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
```

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* 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. 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

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers | **131**

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | **126**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

chap (L2TP)

Syntax

```
chap;
```

Hierarchy Level

```
[edit access group-profile profile-name 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

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266
Configuring an L2TP LNS with Inline Service Interfaces | 261

cisco-nas-port-info (L2TP Tunnel Switching)

Syntax

```
cisco-nas-port-info action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name avp]
```

Release Information
Statement introduced in Junos OS Release 13.2.

Description
Define a tunnel profile for subscriber access.

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.

Options
action—One of the following actions:

- **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.

Default: relay

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 | 175

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**.

chap-secret—(M Series, MX Series, SRX Series, T Series only) For interfaces with PPP encapsulation on which the PPP Challenge Handshake Authentication Protocol (CHAP) is configured, configure the shared secret (the CHAP secret key associated with a peer), as defined in RFC 1994. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *chap-secret*—The secret key associated with a peer.

group-profile—Associate a group profile with a client. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *profile-name*—Name assigned to the group profile.

pap-password—(M Series, MX Series, T Series only) Configure the Password Authentication Protocol (PAP) password. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *password*—PAP password.

user-group-profile—Apply a configured PPP group profile to PPP users. If **user-group-profile** is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.

Values:

- *profile-name*—Name of a PPP group profile configured at the **[edit access group-profile *profile-name*]** hierarchy level.

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

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the CHAP Secret for an L2TP Profile

Configuring PPP CHAP Authentication

Referencing the Group Profile from the L2TP Profile

Configuring the L2TP Client

Configuring Access Profiles for L2TP or PPP Parameters

Configuring the PAP Password for an L2TP Profile

Applying a Configured PPP Group Profile to a Tunnel

Configuring an L2TP Access Profile on the LNS | **268**

delimiter (Access Profile)

Syntax

```
delimiter delimiter;
```

Hierarchy Level

```
[edit access profile profile-name session-options strip-user-name]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

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.

Default

None. You must always configure a delimiter.

Options

delimiter—Character that specifies the boundary between the part of the original username that is kept and the part that is discarded.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

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. 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

- Configuring the L2TP Destination Lockout Timeout | 180
- Configuring an L2TP LNS with Inline Service Interfaces | 261
- Configuring the Reporting and Processing of Subscriber Access Line Information | 247

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
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.

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.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions | 221
- Configuring the L2TP LAC Tunnel Selection Parameters | 219
- LAC Tunnel Selection Overview | 191

destruct-timeout (L2TP)

Syntax

```
destruct-timeout seconds;
```

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.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp destruct-timeout**.

Options

seconds—Length of the destruct timeout.

Range: 10 through 3600

Default: 300

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Setting the L2TP Destruct Timeout | 178

Configuring an L2TP LAC | 183

Configuring an L2TP LNS with Inline Service Interfaces | 261

detection-time

Syntax

```
detection-time {
    threshold 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

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.

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

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

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

device-count (Pseudowire Subscriber Interfaces)

Syntax

```
device-count number;
```

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. The statement also defines the available interface names for the pseudowire interfaces.

> **NOTE:** When you subsequently configure the pseudowire interfaces, you must specify the interface names in the range from ps0 up to ps*(device-count - 1)*. For example, if you set the maximum number of devices to 5, then you can only configure interfaces ps0, ps1, ps2, ps3, and ps4. If you specify an interface name outside that range, the pseudowire interface is not created.

Options

number—Number of devices.

Range: 1 through 7000

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 | 331
- Configuring a Pseudowire Subscriber Logical Interface | 338
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router | 340

dhcp-local-server

Syntax

```
dhcp-local-server {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    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;
            vlan-tags;
        }
    }
    dhcpv6 {
        access-profile profile-name;
        allow-active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
        }
```

```
allow-bulk-leasequery {
    max-connections number-of-connections;
    max-empty-replies number-of-replies;
    restricted-requestor;
    timeout seconds;
}
allow-leasequery {
    restricted-requestor;
}
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;
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                }
                delay-time seconds;
            }
            dual-stack dual-stack-group-name;
            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;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        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;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        dual-stack dual-stack-group-name;
        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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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;
    delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delegated-pool;
    dual-stack dual-stack-group-name;
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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;
            vlan-tags;
        }
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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);
            delay-offer {
                based-on (option-60 | option-77 | option-82) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                }
                delay-time seconds;
            }
            include-option-82 {
                forcerenew;
                nak;
            }
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        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);
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        include-option-82 {
            forcerenew;
            nak;
        }
        dual-stack dual-stack-group-name;
        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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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>;
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    dual-stack dual-stack-group-name;
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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

Understanding Differences Between Legacy DHCP and Extended DHCP

DHCPv6 Local Server Overview

dhcp-relay

Syntax

```
dhcp-relay {
    access-profile profile-name;
    active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
    }
    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;
            vlan-tags;
        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    dhcpv6 {
        access-profile profile-name;
        active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
            topology-discovery;
        }
        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;
        vlan-tags;
    }
}
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;
            vlan-tags;
        }
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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 <strict>;
    use-vlan-id;
}
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>;
    use-vlan-id;
}
```

```
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
  }
  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 <strict>;
    use-vlan-id;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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;
            vlan-tags;
        }
    }
    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 (inet | inet6);
    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 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
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;
            }
        vlan-tags;
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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*
- *Specifying Authentication Support*

dhcpv6 (DHCP Local Server)

Syntax

```
dhcpv6 {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    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;
            vlan-tags;
        }
    }
    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;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

```
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;
    delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
    }
    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;
            vlan-tags;
        }
    }
    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;
            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;
            vlan-tags;
        }
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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 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.

Description

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.

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.

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

dhcp-relay | 489

DHCPv6 Relay Agent Overview

Specifying Authentication Support

dial-options

Syntax

```
dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
}
```

Hierarchy Level

```
[edit interfaces sp-fpc/pic/port unit logical-unit-number],
[edit interfaces si-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces sp-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces si-fpc/pic/port unit logical-unit-number]
```

Release Information

Statement introduced before Junos OS Release 7.4.
The **[edit ...si-...]** hierarchy levels introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level.

l2tp-interface-id *name*—Interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Identifier for Logical Interfaces that Provide L2TP Services

Configuring Dynamic Endpoints for IPsec Tunnels

Configuring Options for the LNS Inline Services Logical Interface | 276

dial-options (Dynamic Profiles)

Syntax

```
dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* **unit** *logical-unit-number*]

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

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level. This option is not currently supported for dynamic profiles.

l2tp-interface-id *name*—(MX Series routers only) L2TP interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Dynamic Profile for Dynamic LNS Sessions | 314

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

Preventing the LAC from Sending Calling Number AVP 22 to the LNS | 252

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 | **321**

drain

Syntax

```
drain;
```

Hierarchy Level

```
[edit services l2tp],
[edit services l2tp destination address ip-address],
[edit services l2tp destination address ip-address routing-instance routing-instance-name],
[edit services l2tp destination name destination-name],
[edit services l2tp tunnel name name],
[edit services l2tp tunnel name name address ip-address],
[edit services l2tp tunnel name name address ip-address routing-instance routing-instance-name]
```

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

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Configuring L2TP Drain | 181
- Configuring an L2TP LAC | 183
- Configuring an L2TP LNS with Inline Service Interfaces | 261

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;
            vlan-tags;
        }
    }
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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

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.

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*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*
- *Configuring RADIUS Reauthentication for DHCP Subscribers*
- *RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers*

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;
            vlan-tags;
        }
    }
    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 (inet | inet6);
    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 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    service-profile dynamic-proflle-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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*

duplicate-clients (DHCPv6 Local Server and Relay Agent)

Syntax

```
duplicate-clients incoming-interface;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server dhcpv6]
```

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Specify the criteria that the **jdhcpd** process uses to support duplicate clients. The router uses the additional criteria to distinguish between the duplicate clients.

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.

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 Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

DHCPv6 Duplicate Client DUIDs | **110**

Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs | **110**

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 | **105**
- Guidelines for Configuring Support for DHCPv4 Duplicate Clients | **105**
- Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces | **108**
- Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information | **106**

dynamic-profile (L2TP)

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Release Information
Statement introduced in Junos OS Release 11.4.

Description
Assign a dynamic profile to the tunnel group for dynamic LNS sessions.

Options
profile-name—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

Configuring a Dynamic Profile for Dynamic LNS Sessions | 314

dynamic-profile (PPP)

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number 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

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | **126**
- *Attaching Dynamic Profiles to MLPPP Bundles*
- For hardware requirements, see *Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces*

dynamic-profiles

Syntax

```
dynamic-profiles {
   profile-name {
      class-of-service {
         dynamic-class-of-service-options {
            vendor-specific-tags tag;
         }
         interfaces {
            interface-name ;
            }
            unit logical-unit-number {
               classifiers {
                  type (classifier-name | default);
               }
               output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
               report-ingress-shaping-rate bps;
               rewrite-rules {
                  dscp (rewrite-name | default);
                  dscp-ipv6 (rewrite-name | default);
                  ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                  inet-precedence (rewrite-name | default);
                  }
               }
            }
         }
         scheduler-maps {
            map-name {
               forwarding-class class-name scheduler scheduler-name;
            }
         }
         schedulers {
            (scheduler-name) {
               buffer-size (seconds | percent percentage | remainder | temporal microseconds);
               drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp |
                  tcp) drop-profile profile-name;
               excess-priority (low | high | $junos-cos-scheduler-excess-priority);
               excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
               overhead-accounting (shaping-mode) <bytes (byte-value>;
               priority priority-level;
               shaping-rate (rate | predefined-variable);
               transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
```

```
        }
    }
    traffic-control-profiles profile-name {
        adjust-minimum rate;
        delay-buffer-rate (percent percentage | rate);
        excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
        excess-rate-high (percent percentage | proportion value);
        excess-rate-low (percent percentage | proportion value);
        guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
        max-burst-size cells;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        peak-rate rate;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate | predefined-variable) <burst-size  bytes>;
        shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
        sustained-rate rate;
    }
}
```

```
firewall {
    family family {
        fast-update-filter filter-name {
            interface-specific;
            match-order [match-order];
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
                only-at-create;
            }
        }
        filter filter-name {
            enhanced-mode-override;
            fast-lookup-filter;
            instance-shared;
            interface-shared;
        interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            only-at-create;
        filter filter-name {
        interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
    hierarchical-policer uid {
        aggregate {
```

```
        if-exceeding {
            bandwidth-limit-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
}
policer uid {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
        policer-action;
    }
}
```

```
    three-color-policer uid {
        action {
            loss-priority high then discard;
        }
        logical-interface-policer;
        single-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            peak-burst-size bytes;
            peak-information-rate bps;
            }
        }
    }
}
```

```
interfaces interface-name {
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
        }
    }
    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;
            }
        }
        encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc |
            atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc
            | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc |
            frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc |
            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);
        family family {
            address address;
            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;
        }
    }
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    service {
        input {
            service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        input-vlan-map {
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (push | swap);
            tag-protocol-id tpid;
            vlan-id number;
        }
        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
        output-vlan-map {
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (pop | swap);
            tag-protocol-id tpid;
            vlan-id number;
        }
        pcef  pcef-profile-name {
            activate rule-name | activate-all;
        }
    }
    unnumbered-address interface-name <preferred-source-address address>;
}
```

```
filter {
    input filter-name (
        shared-name filter-shared-name;
    }
    output filter-name {
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
```

```
        telemetry {
            subscriber-statistics;
            queue-statistics {
                interface $junos-interface-name {
                    refresh rate;
                    queues queue set;
                }
                interface-set $junos-interface-set-name {
                    refresh rate;
                    queues queue set;
                }
            }
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
    }
}
interfaces {
    demux0 {...}
}
interfaces {
    pp0 {...}
}
policy-options {
    prefix-list uid {
        ip-addresses;
        dynamic-db;
    }
}
predefined-variable-defaults predefined-variable <variable-option> default-value;
profile-type remote-device-service;
```

```
protocols {
    igmp {
        interface interface-name {
            accounting;
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave
            log-interval seconds;
            no-accounting;
            oif-map;
            passive;
            promiscuous-mode;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name
            static {
                group group {
                    source source;
                }
            }
            version version;
        }
    }
    mld {
        interface interface-name {
            (accounting | no-accounting);
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave;
            log-interval seconds;
            oif-map;
            passive;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name;
            static {
                group multicast-group-address {
                    exclude;
                    group-count number;
                    group-increment increment;
                    source ip-address {
                        source-count number;
```

```
                    source-increment increment;
                }
            }
        }
        version version;
    }
  }
  router-advertisement {
    interface interface-name {
      current-hop-limit number;
      default-lifetime seconds;
      (managed-configuration | no-managed-configuration);
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
      (other-stateful-configuration | no-other-stateful-configuration);
      prefix prefix;
      reachable-time milliseconds;
      retransmit-timer milliseconds;
    }
  }
}
```

```
routing-instances routing-instance-name {
    interface interface-name;
    routing-options {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        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 {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
    }
}
```

```
routing-options {
    access {
        route prefix {
            next-hop next-hop;
            metric route-cost;
            preference route-distance;
            tag route-tag;
            tag2 route-tag2;
        }
    }
    access-internal {
        route subscriber-ip-address {
            qualified-next-hop underlying-interface {
                mac-address address;
            }
        }
    }
    multicast {
        interface interface-name {
            no-qos-adjust;
        }
    }
}
services {
    captive-portal-content-delivery {
        auto-deactivate value;
        rule name {
            match-direction (input | input-output | output);
            term name {
                then {
                    accept;
                    redirect url;
                    rewrite destination-address address <destination-port port-number>;
                    syslog;
                }
            }
        }
    }
}
```

```
        variables {
            variable-name {
                default-value default-value;
                equals expression;
                mandatory;
                uid;
                uid-reference;
            }
        }
        version-alias profile-alias-string;
    }
}
```

Hierarchy Level

```
[edit]
```

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **filter, policer, hierarchical-policer, three-color-policer, and policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name—Name of the dynamic profile; string of up to 80 alphanumeric characters.

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 a Basic Dynamic Profile*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Dynamic Profiles for Subscriber Management*

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

Enabling the LAC for IPv6 Services | 222
Configuring an L2TP LAC | 183

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.

> **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.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Enabling Tunnel and Global Counters for SNMP Statistics Collection | 163

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

- *Configuring User-Defined Dynamic Variables in Dynamic Profiles*
- *User-Defined Variables*
- *Using Variable Expressions in User-Defined Variables*

equals (Dynamic Profile)

Syntax

```
equals expression;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Configure an expression—a group of arithmetic operators, string operators, and operands—for a user-defined variable that is evaluated at run time and returned as the variable value.

Options

expression—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

- *Configuring User-Defined Dynamic Variables in Dynamic Profiles*
- *User-Defined Variables*
- *Using Variable Expressions in User-Defined Variables*

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
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.

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.

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.

BEST PRACTICE: Use the default method, silent failover.

This statement supersedes the deprecated statement, **disable-failover-protocol**.

Options

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.

silent-failover—Specify silent failover as the resynchronization method.

Default: silent-failover

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 Peer Resynchronization Method | 321

L2TP Failover and Peer Resynchronization | 320

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

Configuring LAC Tunnel Selection Failover Within a Preference Level | 220

Configuring the L2TP LAC Tunnel Selection Parameters | 219

failure-action

Syntax

```
failure-action (clear-binding | clear-binding-if-interface-up | log-only);
```

Hierarchy Level

[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 *group-name* **liveness-detection**],
[edit system services dhcp-local-server dhcpv6 group *group-name* **liveness-detection**],
[edit forwarding-options dhcp-relay group *group-name* **liveness-detection**],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* **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 action the router (or switch) takes when a liveness detection failure occurs.

Options

Default: clear-binding

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.

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.

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.

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*
- *Configuring Detection of DHCP Local Server Client Connectivity with BFD*
- *Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD*
- *Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients*
- *Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients*

flexible-vlan-tagging

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces aex],
[edit interfaces ge-fpc/pic/port],
[edit interfaces et-fpc/pic/port],
[edit interfaces ps0],
[edit interfaces xe-fpc/pic/port]
```

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

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.

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.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Enabling VLAN Tagging*
- *Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers*

forward-snooped-clients (DHCP Local Server)

Syntax

```
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
```

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 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

DHCP Snooping Support

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server | 83

forward-snooped-clients (DHCP Relay Agent)

Syntax

```
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[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 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

Description

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.

The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on non-configured interfaces.

Options

all-interfaces—Perform the action on all interfaces.

Default: On EX Series switches, the action is performed on all interfaces by default.

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 part of a group.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

DHCP Snooping Support

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

fpc (MX Series 5G Universal Routing Platforms)

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. 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

Configuring Port-Mirroring Instances on MX Series 5G Universal Routing Platforms

Enabling Inline Service Interfaces | **273**

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 | 301

gateway-name (Tunnel Profile Remote Gateway)

Syntax

```
gateway-name server-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id 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
server-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 a Tunnel Profile for Subscriber Access | 216

gateway-name (Tunnel Profile Source Gateway)

Syntax

```
gateway-name client-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id 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
client-name—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

Configuring a Tunnel Profile for Subscriber Access | 216

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

Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover | 62

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;
            vlan-tags;
        }
    }
    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;
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
```

```
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
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;
    delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delegated-pool;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

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 Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Understanding Differences Between Legacy DHCP and Extended DHCP*
- *Grouping Interfaces with Common DHCP Configurations*
- *Specifying Authentication Support*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*

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;
            vlan-tags;
        }
    }
    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;
      short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
      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 {
    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 {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
```

```
    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;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

[edit forwarding-options **dhcp-relay**],
[edit forwarding-options dhcp-relay **dhcpv6**],
[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 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

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.

Options

group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.

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*
- *Configuring DHCP Relay Agent*
- *Configuring Group-Specific DHCP Relay Options*
- *Grouping Interfaces with Common DHCP Configurations*
- *Specifying Authentication Support*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 | **266**

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

Syntax

```
hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels number;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 10.1.
implicit-hierarchy option added in Junos OS Release 13.1.
Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.
Support for up to four hierarchy levels added in Junos OS Release 16.1.

Description

Configure hierarchical scheduling options on the interface.

The statement is supported on the following interfaces:

- 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

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.

Options

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.

maximum-hierarchy-levels *number*—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.

- 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.

CAUTION: *MPC3E*, *32x10GE MPC4E*, and *2x100GE + 8x10GE MPC4E* MPCs support only two levels of scheduling hierarchy. When enabling hierarchical scheduling on these cards, you must explicitly set **maximum-hierarchy-levels** to **2**.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

- *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*

holddown-interval

Syntax

```
holddown-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 time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.

Options

milliseconds—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

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

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

hello-interval (L2TP)

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the keepalive timer for L2TP tunnels.

Options

seconds—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of **0** means that no hello messages are sent.

Range: 0 through 3600

Default: 60 seconds

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Timers for L2TP Tunnels

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

identification (Tunnel Profile)

Syntax

```
identification 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 assignment ID of an L2TP tunnel. L2TP sessions with the same tunnel assignment identification and destination are grouped into the same tunnel.

Options

name—Tunnel assignment ID; string of up to 32 alphanumeric characters.

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 | 216

idle-timeout (Access)

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit access group-profile profile-nameppp ppp],
[edit access profile profile-name client client-nameppp ]
```

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:

- 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

Configuring the PPP Attributes for a Group Profile

Configuring PPP Properties for a Client-Specific Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

idle-timeout (L2TP)

Syntax

```
idle-timeout seconds;
```

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.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp tunnel idle-timeout**.

Options

seconds—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.

Range: 0 through 86,400

Default: 60

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Setting the L2TP Tunnel Idle Timeout | **178**

Configuring an L2TP LAC | **183**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

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 Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

For dynamic profiles: Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests | **115**

ignore-magic-number-mismatch (Dynamic Profiles)

Syntax

```
ignore-magic-number-mismatch;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces "$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
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.

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.

For dynamic PPP subscriber connections terminated at the router, configure the statement at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit” ppp-options]** hierarchy level.

For dynamic tunneled PPP subscribers on LNS inline service interfaces, configure the statement at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options]** hierarchy level.

> **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 | **124**

Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests | **115**

For L2TP group profiles: Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

initiate-ncp (Dynamic and Static PPP)

Syntax

```
initiate-ncp (ip | ipv6 | dual-stack-passive);
```

Hierarchy Level

[edit access group-profile *profile-name* ppp **ppp-options**],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" **ppp-options**],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" **ppp-options**],
[edit interfaces pp0 unit *logical-unit-number* **ppp-options**],
[edit interfaces si-*fpc*/*pic*/*port* unit *logical-unit-number* **ppp-options**]

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

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.

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.

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.

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 Network Control Protocol Negotiation Mode | 142

PPP Network Control Protocol Negotiation Mode Overview | 136

inline-services (PIC level)

Syntax

```
inline-services {
   bandwidth (1g | 10g | 20g | 30g | 40g | 100g);
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number]
```

Release Information

Statement introduced in Junos OS Release 11.4.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Enable inline services on PICs residing on MPCs and optionally specify a bandwidth for traffic on the inline service interface.

> **NOTE:** For an MPC, such as MPC2, always configure inline-services at the **[chassis fpc slot-number pic number]** hierarchy level. Do not configure inline services for a service card such as MS-MPC.

The remaining statement is 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

Enabling Inline Service Interfaces | 273
Configuring an L2TP LNS with Inline Service Interfaces | 261

input-hierarchical-policer

Syntax

```
input-hierarchical-policer policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name layer2-policer],
[edit interfaces interface-name unit logical-unit-number 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

policer-name—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

Hierarchical Policers

layer2-policer (Hierarchical Policer)

interface (Dynamic Routing Instances)

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances routing-instance-name]
```

Release Information
Statement introduced in Junos OS Release 9.6.

Description
Assign the specified interface to the dynamically created routing instance.

Options
interface-name—The interface name variable (*$junos-interface-name*). 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 (Service Interfaces)

Syntax

```
interface service-interface-name;
```

Hierarchy Level

```
[edit services service-device-pools pool pool-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.
Support for **ps** service interfaces added in Junos OS Release 19.3R1.

Description

Assign a service interface to a service interface pool. You specify more than one interface for each pool. The interfaces are used to balance traffic loads. For an L2TP tunnel group, the interfaces must be inline service interfaces (si). For a hybrid access tunnel group, the interfaces must be pseudowire service interfaces (ps).

Options

service-interface-name—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

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions | 313
Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278

interface-id

Syntax

```
interface-id interface-id;
```

Hierarchy Level

```
[edit access group-profile profile-name l2tp],
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ike],
[edit access profile profile-name client client-name l2tp],
[edit access profile profile-name client client-name ppp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the interface identifier.

Options

interface-id—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Configuring L2TP for a Group Profile*
- *Configuring the PPP Attributes for a Group Profile*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring an IKE Access Profile*
- Configuring an L2TP Access Profile on the LNS | **268**

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;
    }
}
```

```
        targeted-options {
            backup backup;
            group group;
            primary primary;
            weight ($junos-interface-target-weight | weight-value);
        }
        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;
                local-name name;
            }
            ignore-magic-number-mismatch;
            initiate-ncp (dual-stack-passive | ipv6 | ip)
            ipcp-suggest-dns-option;
            mru size;
            mtu (size | use-lower-layer);
            on-demand-ip-address;
            pap;
            peer-ip-address-optional;
            local-authentication {
                password password;
                username-include {
                    circuit-id;
                    delimiter character;
                    domain-name name;
                    mac-address;
                    remote-id;
                }
            }
        }
        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;
            }
          }
        }
      }
    }
    stacked-interface-set {
      interface-set-name interface-set-name {
        interface-set-name interface-set-name;
      }
    }
}
```

Hierarchy Level

[edit **dynamic-profiles** *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Define interfaces for dynamic client 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*
- *Subscribers over Static Interfaces Configuration Overview*
- *Demultiplexing Interface Overview*

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.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

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. 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

Configuring IP Inline Reassembly for L2TP | 318

IP Packet Fragment Reassembly for L2TP Overview | 317

ip-reassembly (L2TP)

Syntax

```
ip-reassembly {
    service-set service-set-name;
}
```

Hierarchy Level

[edit services **l2tp**]

Release Information

Statement introduced in Junos OS Release 13.1.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Associate the reassembly service-set with the L2TP service.

> **NOTE:** The service set must be defined at the **[edit services]** hierarchy level.

Options

service-set *service-set-name*—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

IP Packet Fragment Reassembly for L2TP Overview | 317
Configuring IP Inline Reassembly for L2TP | 318

ip-reassembly-rules (Service Set)

Syntax

```
ip-reassembly-rules rule-name;
```

Hierarchy Level

```
[edit services service-setservice-set-name]
```

Release Information

Statement introduced in Junos OS Release 13.1.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

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

rule-name—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

Configuring IP Inline Reassembly for L2TP | **318**

IP Packet Fragment Reassembly for L2TP Overview | **317**

ipcp-suggest-dns-option

Syntax

```
ipcp-suggest-dns-option;
```

Hierarchy Level

[edit access group-profile *group-profile-name* ppp **ppp-options**],
[edit dynamic-profiles *profile-name* interfaces pp0 **unit** "$junos-interface-unit" **ppp-options**],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" **unit** "$junos-interface-unit" **ppp-options**],
[edit interfaces pp0 unit *logical-unit-number* **ppp-options**],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number* **ppp-options**]

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

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses | **144**

Configuring the PPP Attributes for a Group Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Configuring Dynamic Authentication for PPP Subscribers | **131**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

keepalive

Syntax

```
keepalive seconds;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the keepalive interval for an L2TP tunnel.

Options

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.

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.

Range: 0 through 32,767 seconds

Default: 30 seconds

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the PPP Attributes for a Group Profile

Configuring PPP Properties for a Client-Specific Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

keepalives

Syntax

```
keepalives <interval seconds> <down-count number> <up-count number>;
```

Hierarchy Level

```
[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]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.

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:

- **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

down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes down a link.

Range: 1 through 255

Default: 3

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 10 seconds

up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up.

Range: 1 through 255

Default: 1

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Keepalives

Configuring Frame Relay Keepalives

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

keepalives (Dynamic Profiles)

Syntax

```
keepalives {
    interval seconds;
    }
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit logical-unit-number ]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]
```

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 10.1.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$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.

Starting in Junos OS Release 15.1R5, you can configure the PPP keepalive interval for subscriber services in the range 1 second through 600 seconds. Subscriber PPP keepalives are handled by the Packet Forwarding Engine. If you configure a value greater than 600 seconds, the number is accepted by the CLI, but the Packet Forwarding Engine limits the interval to 600 seconds.

In earlier Junos OS releases, the range is from 1 second through 60 seconds. The Packet Forwarding Engine limits any higher configured value to an interval of 60 seconds.

PPP keepalives for nonsubscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

Default

Sending of keepalives is enabled by default.

Options

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 600 seconds for subscriber services

Range: 1 through 32767 seconds for nonsubscriber services

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

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers | **131**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

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);
    failover-resync (failover-protocol | silent-failover);
    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.

> **NOTE:** L2TP is not supported over GRE tunnels.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **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 | **153**

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.

> **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.

Options

interface-id—Configure the interface identifier.

Values:

- *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

lcp-renegotiation—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.

NOTE: This statement is not supported at the **[edit access group-profile l2tp]** hierarchy level for L2TP LNS on MX Series routers.

local-chap—Configure the Junos OS so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

maximum-sessions—Specify the maximum number of L2TP sessions for the chassis, all tunnels, a tunnel group, a session limit group, or a client.

Values:

- *number*—Number of sessions allowed.
- Range: (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit
- Range: (Tunnel) 1 through 65,536

maximum-sessions-per-tunnel—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.

Values:

- *number*—Maximum number of sessions for a Layer 2 tunnel.

multilink—Configure Multilink PPP for Layer 2 Tunneling Protocol (L2TP).

The options for this statement are explained separately. Click the linked statement for details.

override-result-code—Configure the LNS to override result codes in Call-Disconnect-Notify (CDN) messages.

Values:

- 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.

ppp-authentication—(T Series only) Configure PPP authentication.

> **NOTE:** This statement is not supported for L2TP LNS on MX Series routers.

Values:

- chap—Challenge Handshake Authentication Protocol.
- pap—Password Authentication Protocol.

ppp-profile—(M Series, T Series only) Specify the profile used to validate PPP session requests through L2TP tunnels.

> **NOTE:** This statement is not supported for L2TP LNS on MX Series routers.

Values: *profile-name*—Identifier for the PPP profile.

sessions-limit-group—(MX Series only) Starting in Junos OS Release 16.1, specify in an L2TP access profile the session limit group to which a client is assigned by the profile.

Values: *limit-group-name*—Identifier of the session-limit group to which a client is assigned.

service-profile—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.

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.

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.

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.

Values:

- *profile-name*—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 (&).
- *parameter*—(Optional) Value to be passed to the service when it is activated on the subscriber session.

shared-secret—Configure the shared secret.

Values:

- *shared-secret*—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

Configuring L2TP Properties for a Client-Specific Profile

Configuring an L2TP Access Profile on the LNS | **268**

Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | **213**

Configuring an L2TP LAC | **183**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

L2TP for Subscriber Access Overview | **153**

l2tp-access-profile

Syntax

```
l2tp-access-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the profile used to validate all L2TP connection requests to the local gateway address.

Options

profile-name—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

Configuring Access Profiles for L2TP Tunnel Groups

Configuring an L2TP Access Profile on the LNS | **268**

l2tp-maximum-session (Service Interfaces)

Syntax

```
l2tp-maximum-session number;
```

Hierarchy Level

```
[edit interfaces si-slot/pic/port],
[edit interfaces asinumber]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

Specify the maximum number of L2TP sessions allowed on a physical service interface (si) or aggregated service interface (asi).

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.

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.

Options

number—Maximum number of L2TP sessions allowed for the interface. A value of 0 prevents the interface from being considered.

Default: 64,000

Range: 0 through 64,000

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | **213**

L2TP Session Limits and Load Balancing for Service Interfaces | **283**

Configuring an L2TP LAC | **183**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

L2TP for Subscriber Access Overview | **153**

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
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.

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 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.

This behavior can be referred to as the *receive* functionality for BNG Layer 2 liveness detection, as opposed to the *send* functionality configured with the **layer2-liveness-detection (Send)** statement for DHCP relay or DHCP local server.

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.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*
- *DHCP Liveness Detection Overview*
- *Configuring Junos OS Enhanced Subscriber Management*

layer2-liveness-detection (Send)

Syntax

```
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval seconds;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection **method**],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection **method**],
[edit forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name* liveness-detection **method**],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection **method**],
[edit forwarding-options dhcp-relay liveness-detection **method**],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection **method**],
[edit system services dhcp-local-server dhcpv6 liveness-detection **method**],
[edit system services dhcp-local-server **dual-stack-group** *dual-stack-group-name* liveness-detection **method**],
[edit system services dhcp-local-server group *group-name* liveness-detection **method**],
[edit system services dhcp-local-server liveness-detection **method**],

Release Information

Statement introduced in Junos OS Release 17.4R1.

Description

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.

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.

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.

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.

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.

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.

> **NOTE:** The only option to the **failure-action** statement supported by Layer 2 liveness detection is **clear-binding**.

Options

max-consecutive-retries *number*—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

transmit-interval *seconds*—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

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 Using ARP and Neighbor Discovery Packets*
- *DHCP Liveness Detection Overview*

lcp-renegotiation

Syntax

```
lcp-renegotiation;
```

Hierarchy Level

```
[edit access group-profile profile-name l2tp],
[edit access profile profile-name client client-name 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.

> **NOTE:** This statement is not supported at the **[edit access group-profile l2tp]** hierarchy level 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 for a Group Profile*
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS | **268**

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*
- *Configuring Detection of DHCP Local Server Client Connectivity with BFD*
- *Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD*
- *Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients*
- *Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*

local-authentication (Dynamic PPP Options)

Syntax

```
local-authentication {
    password password;
    username-include {
        circuit-id;
        delimiter character;
        domain-name name;
        mac-address;
        remote-id;
    }
}
```

Hierarchy Level

[edit dynamic-profiles *name* interfaces $junos-interface-ifd-name **unit** $junos-interface-unit **ppp-options**]

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Configure local authentication for terminated PPP subscribers. This enables the external RADIUS server to pass implementation-specific configuration for successfully authenticated subscribers. Local authentication enables the same dynamic profile to support both CPEs that do not negotiate authentication protocols and CPEs that use PAP or CHAP authentication.

Options

password *password*—Specify the local authentication password

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface

RELATED DOCUMENTATION

Configuring Local Authentication in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | 128

local-gateway (L2TP LNS)

Syntax

```
local-gateway {
    address address;
    gateway-name gateway-name;
}
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

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.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Local Gateway Address and PIC.

Configuring L2TP Tunnel Groups

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301

lockout-timeout (L2TP Destination Lockout)

Syntax

```
lockout-timeout seconds;
```

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.

> **NOTE:** The *ip-address* option for the **destination** statement does not apply to the **lockout-timeout** statement.

Options

seconds—Length of the period during which the destination is locked out.

Range: 60 through 3600 seconds

Default: 300 seconds

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 | **180**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

logical-system (Tunnel Profile)

Syntax

```
logical-system logical-system-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

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

logical-system-name— Name of the logical system.

Default: Logical system *default*

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 | 216

mac

Syntax

```
mac mac-address;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

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
mac-address—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. 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

Configuring the MAC Address on the Management Ethernet Interface

Configuring a Pseudowire Subscriber Logical Interface Device | 341

mac-address (Dynamic Access-Internal Routes)

Syntax

```
mac-address address;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options access-internal route
    subscriber-ip-address qualified-next-hop underlying-interface],
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name
    access-internal route subscriber-ip-address qualified-next-hop underlying-interface],
[edit dynamic-profiles routing-options access-internal route subscriber-ip-address qualified-next-hop
    underlying-interface]
```

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 route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the MAC address variable for an access-internal route for unnumbered interfaces such as DHCP subscriber interfaces.

Options

address—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

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers | 66

match-direction (IP Reassembly Rule)

Syntax

```
match-direction direction
```

Hierarchy Level

```
[edit services ip-reassemblyrulerule-name]
```

Release Information

Statement introduced in Junos OS Release 13.1.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

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

direction—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

Configuring IP Inline Reassembly for L2TP | **318**
IP Packet Fragment Reassembly for L2TP Overview | **317**

maximum-sessions (L2TP)

Syntax

```
maximum-sessions number;
```

Hierarchy Level

```
[edit access profile profile-name client client-name],
[edit services l2tp],
[edit services l2tp sessions-limit-group],
[edit services l2tp tunnel],
[edit services l2tp tunnel-group group-name],
```

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

number—Number of sessions allowed.

Range: (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit

Range: (Tunnel) 1 through 65,536

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | **213**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**
- L2TP for Subscriber Access Overview | **153**

maximum-sessions-per-tunnel

Syntax

```
maximum-sessions-per-tunnel number;
```

Hierarchy Level

```
[edit access group-profile l2tp],
[edit access profile profile-name client client-name l2tp]
```

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

number—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

- *Configuring L2TP for a Group Profile*
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS | **268**

max-sessions (Tunnel Profile)

Syntax

```
max-sessions number;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the maximum number of sessions allowed in the tunnel.

Options

number—Maximum number of sessions allowed in the tunnel. A value of 0 means that the maximum configurable number of sessions is allowed.

Range: 0 through 60,000

Default: 0

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 | 216

medium (Tunnel Profile)

Syntax

```
medium type;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Specify the medium type for the tunnel.

Default
ipv4

Options
type—Medium type for the tunnel. The only value currently available is **ipv4**.

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 | 216

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 Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *DHCP Liveness Detection Overview*
- *Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients*
- *Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*

metric (Dynamic Access-Internal Routes)

Syntax

```
metric route-cost;
```

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 the **[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 cost for an access route.

Options

route-cost—Either the specific cost you want to assign to the access route or either of the following cost variables:

- **$junos-framed-route-cost**—Cost of an IPv4 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-cost**—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

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 | 64

minimum-interval

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd transmit-interval]
```

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 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 **transmit-interval minimal-interval** and **minimum-receive-interval** statements.

Options

milliseconds — Specify the minimum interval value for BFD liveliness detection.

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

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

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

minimum-receive-interval

Syntax

```
minimum-receive-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 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

milliseconds — 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

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

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

minimum-retransmission-timeout (L2TP Tunnel)

Syntax

```
minimum-retransmission-timeout seconds;
```

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

seconds—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

- Configuring Retransmission Attributes for L2TP Control Messages | 160
- Configuring an L2TP LAC | 183
- Configuring an L2TP LNS with Inline Service Interfaces | 261

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 Metro 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

```
multiplier number;
```

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 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

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

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

name (L2TP Destination)

Syntax

```
name destination-name {
    drain;
}
```

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

destination-name—Locally assigned name of the tunnel destination.

The remaining statement is 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

Configuring L2TP Drain | 181

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. 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

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 *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 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

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

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

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.

> **NOTE:** This global configuration can be overridden by the configuration in a tunnel profile or by the RADIUS configuration.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Globally Configuring the LAC to Interoperate with Cisco LNS Devices | 189
- Configuring a Tunnel Profile for Subscriber Access | 216
- LAC Interoperation with Third-Party LNS Devices | 189

nas-port-method (Tunnel Profile)

Syntax

```
nas-port-method cisco-avp;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

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

Configuring a Tunnel Profile for Subscriber Access | 216

next-hop (Dynamic Access Routes)

Syntax

```
next-hop next-hop;
```

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 the **[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 next-hop address for an access route. Access routes are typically unnumbered interfaces.

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).

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.

Options

next-hop—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.

- 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

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 | 64

next-hop-service

Syntax

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

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

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

inside-service-interface *interface-name.unit-number*—Name and logical unit number of the service interface associated with the service set applied inside the network.

outside-service-interface *interface-name.unit-number*—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type *interface-type*—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.

service-interface-pool *name*—Name of the pool of logical interfaces configured at the **[edit services service-interface-pools pool *pool-name*]** 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.

> **NOTE:** **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

no-allow-snooped-clients

Syntax

```
no-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.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Explicitly disable DHCP snooping support on DHCP relay agent.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to explicitly disable snooping support for DHCPv6 relay agent.

NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is *enabled* by default. In Release 10.1 and later, DHCP snooping is *disabled* by default.

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

Overriding the Default DHCP Relay Configuration Settings

DHCP Snooping Support

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

no-gratuitous-arp-request

Syntax

```
no-gratuitous-arp-request;
```

Hierarchy Level

[edit interfaces *interface-name*]

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 Metro 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

Configuring Gratuitous ARP

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 logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6]
```

Release Information

Statement introduced in Junos OS Release 15.1R2.

Description

Disable 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 if 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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Disabling DHCP Snooping Filters | 93

on-demand-ip-address

Syntax

```
on-demand-ip-address;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”].
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit interfaces pp0 unit unit-number ppp-options], [edit protocols ppp-service]
```

Release Information

Statement introduced in Junos OS Release 13.1.

Description

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.

Configuration changes take effect as follows:

- 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.

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.

Default

This functionality is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

options (Access Profile)

Syntax

```
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;
        nas-identifier;
    }
    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;
    }
    ip-address-change-notify message;
    juniper-access-line-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
            port-width width;
            pw-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;
        }
    }
    override {
        calling-station-id remote-circuit-id;
        nas-ip-address tunnel-client-gateway-address;
        nas-port tunnel-client-nas-port;
        nas-port-type tunnel-client-nas-port-type;
    }
    remote-circuit-id-delimiter;
    remote-circuit-id-fallback;
    remote-circuit-id-format {
        agent-circuit-id;
        agent-remote-id;
    }
    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;
}
```

Hierarchy Level

[edit access profile *profile-name* radius]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
juniper-dsl-attributes introduced in Junos OS Release 11.4.
nas-port-id-delimiter introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
calling-station-id-delimiter introduced in Junos OS Release 13.1.
ip-address-change-notify introduced in Junos OS Release 13.1.
coa-dynamic-variable-validation, **client-authentication-algorithm**, and **client-accounting-algorithm** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
remote-circuit-id-delimiter, **remote-circuit-id-fallback**, and **remote-circuit-id-format** introduced in Junos OS Release 13.3R1 on MX Series.
chap-challenge-in-request-authenticator introduced in Junos OS Release 15.1.
nas-identifier introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.
service-activation introduced in Junos OS Release 16.2.
juniper-access-line-attributes introduced in Junos OS Release 19.2R1

Description

Configure the options used by RADIUS authentication and accounting servers.

Options

accounting-session-id-format—(EX Series, MX Series only) Configure the format the router or switch uses to identify the accounting session. The default is **decimal**.

Values:

- decimal—Use the decimal format.
- description—Use the generic format, in the form: **jnpr *interface-specifier*:*subscriber-session-id***.

calling-station-id-delimiter—(MX Series, T Series only) Starting in Junos OS Release 13.1, specify the character that the router uses as a separator between the concatenated values in the Calling-Station-ID (RADIUS IETF attribute 31) string. The router uses the delimiter when you configure more than one value in the **calling-station-id-format** statement. The default is the hash (**#**) character.

Values:

- *delimiter-character*—Character to use for the delimiter. You must enclose the delimiter character in quotation marks (“ ”).

chap-challenge-in-request-authenticator—(MX Series only) Starting in Junos OS Release 15.1, configure the **authd** process to insert the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets, if the challenge value is 16 bytes long. If you enable the **chap-challenge-in -request-authenticator** statement and the random challenge is not 16 bytes long, **authd** ignores the statement and uses the default behavior, which inserts the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets.

client-accounting-algorithm—(EX Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the access method the router uses to access RADIUS accounting servers. The default is the **direct** option.

Values:

- direct—Use the direct method.
- round-robin—Use the round-robin method.

client-authentication-algorithm—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the method that the authenticator uses to access RADIUS authentication servers when there are multiple servers configured. Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The router or switch, acting as the authenticator, waits for a response from the server before sending another request.

When there are multiple RADIUS server connections configured for a client, the authenticator attempts to reach the different servers in the order that they are configured. If there is no response from the first RADIUS server, the authenticator attempts to reach the next RADIUS server. This process repeats until the client is either granted access or there are no more configured servers.

If the **direct** method is configured, the authenticator always treats the first server in the list as the primary server. The authenticator moves on to the second server only if the attempt to reach the first

server fails. If the **round-robin** method is configured, the server chosen first will be rotated based on which server was used last. The first server in the list is treated as a primary for the first authentication request, but for the second request, the second server configured is treated as primary, and so on. With this method, all of the configured servers receive roughly the same number of requests on average so that no single server has to handle all of the requests.

NOTE: The **round-robin** access method is not recommended for use with EX Series switches.

Default: The default is the **direct** option.

Values:

- direct—Use the direct access method. The authenticator contacts the first RADIUS server on the list for each request, the second server if the first one fails, and so on.
- round-robin—Use the round-robin method. The authenticator contacts the first RADIUS server for the first request, the second server for the second request, and so on.

coa-dynamic-variable-validation—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, specify that when a CoA operation includes a change to a client profile dynamic variable that cannot be applied (such as an update to a non-existent filter), the router does not apply any changes to client profile dynamic variables in the request, and responds with a NACK message.

Default: If you do not configure this statement, the router does not apply any incorrect variable updates, but does make any other changes to the client profile dynamic variables, and responds with an ACK message.

ethernet-port-type-virtual—(EX Series, M Series, MX Series only) Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of **ethernet** in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of **virtual**.

NOTE: This statement takes precedence over the **nas-port-type** statement if you include both statements in the same access profile.

access-loop-id-local—Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

ip-address-change-notify—(MX Series only) Starting in Junos OS Release 13.1, 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. The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

Optionally, configure a message that is included in the VSA when it is sent to the RADIUS server.

Default: This functionality is disabled by default.

Values: ***message***—VSA message.

Range: Up to 32 characters.

juniper-access-line-attributes—Configure AAA to add Juniper Networks access line VSAs to the RADIUS authentication and accounting request messages for subscribers. If the router has not received and processed the corresponding ANCP attributes from the access node, then AAA provides only the following in these RADIUS messages:

- Downstream-Calculated-QoS-Rate (IANA 4874, 26-141)—Default configured advisory transmit speed.
- Upstream-Calculated-QoS-Rate (IANA 4874, 26-142)—Default configured advisory receive speed.

NOTE: Starting in Junos OS Release 19.2R1, the **juniper-access-line-attributes** option replaces the **juniper-dsl-attributes** option. The difference between these options is that **juniper-dsl-attributes** supported only DSL TLVs received in the ANCP Port Status message. The **juniper-access-line-attributes** option supports PON TLVs in addition to DSL TLVs, and will be extensible to future access technologies.

For backward compatibility with existing scripts, the **juniper-dsl-attributes** option redirects to the new **juniper-access-line-attributes** option. We recommend that you use **juniper-access-line-attributes**.

NOTE: The **juniper-access-line-attributes** option is not backward compatible with Junos OS Release 19.1 or earlier releases. This means that if you have configured **juniper-access-line-attributes** option in Junos OS Release 19.2 or higher releases, you must perform the following steps to downgrade to Junos OS Release 19.1 or earlier releases:

1. Delete the **juniper-access-line-attributes** option from all access profiles that include it.
2. Perform the software downgrade.
3. Add the **juniper-dsl-attributes** option to the affected access profiles.

Default: The Juniper Networks access line VSAs are not added to the RADIUS authentication and accounting request messages. However, the DSL Forum VSA—if available—is added to RADIUS messages by default.

nas-identifier—(EX Series, MX Series, SRX Series only) Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests. This statement was introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

Values: *identifier-value*—String to use for authentication and accounting requests.

Range: 1 through 64 characters.

nas-port-id-delimiter—(MX Series only) Starting in Junos OS Release 11.4, specify the character that the router uses as a separator between the concatenated values in the NAS-Port-ID string. The router uses the delimiter when you configure more than one value in the **nas-port-id-format** statement. The default is the hash (**#**) character. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Values: *delimiter-character*—Character used for the delimiter.

remote-circuit-id-delimiter—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure a delimiter character for the remote circuit ID string when you use the **remote-circuit-id-format** statement to configure the string to use instead of the Calling-Station ID in L2TP Calling Number AVP 22. If more than one value is configured for the remote circuit ID format, the delimiter character is used as a separator between the concatenated values in the resulting remote circuit ID string. The default is the hash (**#**) character.

Values: *delimiter*—Delimiter character to be used between components of the remote circuit ID string.

remote-circuit-id-fallback—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the fallback value for the LAC to send in L2TP Calling Number AVP 22, either the configured Calling-Station-ID or the default underlying interface. Use of the fallback value is triggered when the components of the override string you configured with the **remote-circuit-id-format** statement—the ACI, the ARI, or both ACI and ARI—are not received by the LAC in the PPPoE Active Discovery Request (PADR) packet.

Values:

- configured-calling-station-id—Send the configured Calling-Station-ID in the Calling Number AVP.
- default—Send the underlying interface value in the Calling Number AVP.

remote-circuit-id-format—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the format of the string that overrides the Calling-Station-ID format in the Calling Number AVP 22 sent by the LAC to the LNS in the ICRQ packet when an L2TP session is being established. You can specify the ACI, the ARI, or both the ACI and ARI. This statement enables you to decouple the AVP 22 value from the RADIUS Calling-Station-ID attribute (31); the values for AVP 22 and the Calling-Station-ID attribute are the same when you use the **calling-station-id-format** statement to configure AVP 22.

NOTE: You must configure the **override calling-circuit-id remote-circuit-id** statement for the remote circuit ID format to be used in the calling number AVP.

Values:

- agent-circuit-id—Specifies use of the ACI string that uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. For PPPoE traffic, the ACI string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.
- agent-remote-id—Specifies use of the ARI string that identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The agent remote identifier (ARI) string is stored in the DSL Forum Agent-Remote-ID VSA [26-2] for PPPoE traffic.

service-activation—(MX Series only) Starting in Junos OS Release 16.2, specify whether subscribers are allowed to log in even when service activation failures related to configuration errors occur during family activation request processing by authd for a newly authenticated subscriber. Configuration errors include missing or incorrect syntax, missing or incomplete references to dynamic profiles, and missing or incomplete variables.

NOTE: This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policies.

You can enable separate configurations for subscriber login services for two **service-activation** types: **dynamic-profile** and **extensible-service**. You configure the **dynamic-profile** type services in the dynamic profile at the **[edit dynamic-profiles]** hierarchy level; the profile is used to provide dynamic subscriber access and services for broadband applications. The **extensible-service** type is for business services configured in an operation script and provisioned by the Extensible Subscriber Services Manager daemon (essmd).

Default:

Default behavior depends on the service type:

- For **extensible-service** services: **optional-at-login**.
- For **dynamic-profile** services: **required-at-login**.

Values:

- optional-at-login—Service activation is optional. Failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Failure for any other reason causes network family activation to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.
- required-at-login—Service activation is required. Failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.

vlan-nas-port-stacked-format—(MX Series only) Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

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

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Configuring Access Profile Options for Interactions with RADIUS Servers*
- *RADIUS Servers and Parameters for Subscriber Access*
- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Configuring a Calling-Station-ID with Additional Options*

override (RADIUS Options)

Syntax

```
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
```

Hierarchy Level

[edit access profile *profile-name* radius **options**]

Release Information

Statement introduced in Junos OS Release 13.3R1 on MX Series.
nas-ip-address tunnel-client-gateway-address and **nas-port tunnel-client-nas-port** options added in Junos OS Release 17.3R1 on MX Series.
calling-station-id remote-circuit-id option added in Junos OS Release 13.3R1 on MX Series.

Description

Override the values for certain RADIUS options.

Options

calling-station-id remote-circuit-id—Override the Calling-Station-ID format for the string that the LAC sends to the LNS in the Calling Number AVP 22. AVP 22 is included in the incoming-call request (ICRQ) packet when the L2TP session is being established.

If you have configured the **calling-station-id-format** statement, the LAC sends AVP 22 in a different format that can include other information in addition to the ACI or ARI. In this case, the LNS sends the value of AVP 22 to the RADIUS server as the Calling-Station-ID attribute (31).

If you do not want the AVP 22 value to be used as the Calling-Station-ID attribute (31) for some subscribers, you can override that format in the access profile by including the **calling-station-id remote-circuit-id** override option. You must also configure the new format to be used with the **remote-circuit-id-format** statement; you can specify: the ACI, ARI, or both the ACI and ARI received from the PADR packet.

nas-ip-address tunnel-client-gateway-address—Check the SDB to determine whether the session's LAC endpoint IP address is available. If so, use that value in the RADIUS NAS-IP-Address attribute (4). The LNS subsequently sends the attribute to the RADIUS server in Access-Request and accounting messages. If the value is not present, the attribute is not sent.

nas-port tunnel-client-nas-port—Check the SDB to determine whether the LAC NAS port information is conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If so, use that value in the RADIUS NAS-Port attribute (5). The LNS subsequently sends the attribute to the RADIUS server in Access-Request and accounting messages. If the value is not present, the LNS sends its local IP address.

nas-port-type tunnel-client-nas-port-type—Check the SDB to determine whether the LAC NAS port information is conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If so, override the value of the RADIUS NAS-Port-Type attribute (61) at the LNS with that value. Otherwise, use the original NAS-Port-Type in attribute 61.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Override the Calling-Station-ID Format for the Calling Number AVP | 253

RADIUS Servers and Parameters for Subscriber Access

overrides (DHCP Relay Agent)

Syntax

```
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;
}
```

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 dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[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 8.3.

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

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.

M120 and M320 routers do not support DHCPv6.

The following statements are supported at both the **[edit ... dhcp-relay]** and **[edit ... dhcpv6]** hierarchy levels.

- **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. 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*
>
> *Overriding the Default DHCP Relay Configuration Settings*

overrides (Enhanced Subscriber Management)

Syntax

```
overrides {
    event {
        catastrophic-failure {
            reboot (master | standby);
        }
    }
    interfaces {
        family (inet | inet6) {
            layer2-liveness-detection;
            ipoe-dynamic-arp-enable;
            receive-gratuitous-arp;
        }
    }
    no-unsolicited-ra;
    ra-initial-interval-max seconds;
    ra-initial-interval-min seconds;
    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.

ra-initial-interval-max and **ra-initial-interval-min** options added in Junos OS Release 18.2R1 on MX Series routers.
ipoe-dynamic-arp-enable and **receive-gratuitous-arp** options added in Junos OS Release 18.4R1 on MX Series routers.

Description

Override the default configuration settings for the Junos OS enhanced subscriber management software for subscriber management.

Options

ra-initial-interval-max *seconds*—Specify the high end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the **ra-initial-interval-min** option.

Range: 1 through 16

ra-initial-interval-min *seconds*—Specify the low end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the **ra-initial-interval-max** option.

BEST PRACTICE: Always configure the value of **ra-initial-interval-min** to be less than or equal to the value of **ra-initial-interval-max**. If you configure the values to be the same, the initial router advertisement intervals are constant and not randomized.

Range: 1 through 16

ipoe-dynamic-arp-enable—Enable dynamic ARP to resolve the MAC address for IPv4 framed host (32-bit) routes. By default the framed route is permanently associated with the source MAC address received in the packet that triggered creation of the dynamic VLAN.

receive-gratuitous-arp—Enable the router to compare the source MAC address received in a gratuitous ARP request or reply packet with the value in the ARP cache. The router updates the cache with the received MAC address when it determines this address is different from the cache entry.

This situation occurs when an IPv4 address is moved to a different device. The device broadcasts a gratuitous ARP reply packet with its MAC address as the source MAC address. When the **receive-gratuitous-arp** option is configured, the router compares the MAC addresses and updates the cache to associate the IPv4 address with the new MAC address.

If the **receive-gratuitous-arp** option is not configured, the router does not accept the gratuitous ARP request or reply packet and cannot quickly learn about the new address. Instead, the original dynamic ARP entry in the cache eventually times out. Before deleting the entry, the router sends an ARP request for the target IP address. The client responds with the new MAC address. This delay in learning about the new address means there is a period during which the MAC address in the ARP cache does not match the address in the new device's NIC.

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*
- *Junos OS Enhanced Subscriber Management Overview*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*
- *Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors*

pap

Syntax

```
pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* **ppp-options**],
[edit interfaces *interface-name* unit *logical-unit-number* **ppp-options**],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **ppp-options**]

Release Information

Statement introduced in Junos OS Release 8.3.

Description

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.

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.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **pap** statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.

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

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol On a Logical Interface*
- *Tracing Operations of the pppd Process*
- *traceoptions (PPP Process)*
- *Example: Configuring PAP for an L2TP Profile*
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

pap (Dynamic PPP)

Syntax

```
pap;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit” **ppp-options**],
[edit dynamic-profiles *profile-name* interfaces "$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 *profile-name* 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

- *Dynamic Profiles Overview*
- Configuring Dynamic Authentication for PPP Subscribers | **131**
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | **126**
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

pap (L2TP)

Syntax

```
pap;
```

Hierarchy Level

```
[edit access group-profile profile-name 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

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

parse-direction (Access Profile)

Syntax

```
parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access profile profile-name session-options strip-user-name]
```

Release Information
Statement introduced in Junos OS Release 16.2.

Description
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. characters become the new, modified username.

Default
left-to-right

Options
left-to-right—Parse the subscriber login string from left to right up to the delimiter.

For example, when the direction is **left-to-right**, the characters **/@$%#** are configured as the delimiters, and the login string is **drgt21@example.com$84**, the **@** is reached before the **$**, so the username is modified to **drgt21**.

right-to-left—Parse the subscriber login string from right to left up to the delimiter.

For example, when the direction is **right-to-left**, the characters **/@$%#** are configured as the delimiters, and the login string is **drgt21@example.com$84**, the **$** is reached before the **@**, so the username is modified to **drgt21@example.com**.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

pic (M Series and T Series Routers)

Syntax

```
pic pic-number {
    ce1 {
        e1 port-number {
            channel-group group-number timeslots slot-number;
        }
    }
    ct3 {
        port port-number {
            t1 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. 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

- *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 (Service Interfaces)

Syntax

```
pool pool-name {
    interface service-interface-name;
}
```

Hierarchy Level

```
[edit services service-device-pools]
```

Release Information

Statement introduced in Junos OS Release 11.4.
Support for **ps** service interfaces added in Junos OS Release 19.3R1.

Description

Define a pool of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The interfaces in the pool must be inline service interfaces (si). The service device pool is required for dynamic LNS sessions.

Options

pool-name—Name of the service interface pool.

The remaining statement is 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

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions | 313

pp0 (Dynamic PPPoE)

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. 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

Configuring a PPPoE Dynamic Profile

Configuring Dynamic Authentication for PPP Subscribers | **131**

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.
Statement **idle-timeout** introduced in Junos OS Release 11.1 for the QFX Series.
Statement **idle-timeout** introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description

Configure PPP properties for a group profile.

Options

cell-overhead—(M Series, MX Series, PTX Series, T Series only) Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

encapsulation-overhead—(MX Series, T Series only) Configure the encapsulation overhead for class-of-service calculations.

Values: *bytes*—The number of bytes used as encapsulation overhead for the session.

framed-pool—(M Series, MX Series, PTX Series, T Series only) Configure the address pool.

Values: *framed-pool*—References a configured address pool.

idle-timeout—(EX4600, M Series, MX Series, OCX1100, PTX Series, QFX Series, T Series only) Configure the idle timeout for a user. Starting in Junos OS Release 11.1, this statement is available on the QFX Series. Starting in Junos OS Release 14.1X53-D20, this statement is available on OCX Series switches.The router might consider a PPP session to be idle because of the following reasons:

- 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.

Values: *seconds*—Number of seconds a user can remain idle before the session is terminated.

Range: 0 through 4,294,967,295 seconds

Default: 0

interface-id *interface-id*—(M Series, MX Series, PTX Series, T Series only) Configure the interface identifier.

Values: *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

keepalive—(M Series, MX Series, PTX Series, T Series only) Configure the keepalive interval for an L2TP tunnel.

Values: *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.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.

Range: 0 through 32,767 seconds

Default: 30 seconds

primary-dns—(EX Series, SRX Series only) Configure the primary Domain Name System (DNS) server.

Values: *primary-dns*—An IPv4 address.

primary-wins—(M Series, MX Series, PTX Series, T Series only) Configure the primary Windows Internet name server.

Values: *primary-wins*—An IPv4 address.

secondary-dns—(SRX Series only) Configure the secondary DNS server.

Values: *secondary-dns*—An IPv4 address.

secondary-wins—(M Series, MX Series, PTX Series, SRX Series, T Series only) Configure the secondary Windows Internet name server.

Values: *secondary-wins*—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the PPP Attributes for a Group Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**

Configuring PPP Properties for a Client-Specific Profile

ppp-options

Syntax

```
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;
    }
}
```

Hierarchy Level

```
[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]
```

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. 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

Configuring the PPP Challenge Handshake Authentication Protocol

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

ppp-options (Dynamic PPP)

Syntax

```
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    lcp-connection-update;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”].
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
```

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.
lcp-connection-update option added in Junos OS Release 20.2R1.

Description

Configure PPP-specific interface properties in a dynamic profile.

> **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.

Options

lcp-connection-update—Enable PPP to act on a Connection-Status-Message VSA (26–218) received by authd in either a RADIUS Access-Accept message or a CoA message. PPP conveys the contents of the VSA in an LCP Connection-Update-Request message to the remote peer, such as a home gateway. This action requires the following to be true:

- At least the first address family has been successfully negotiated and the session is active.
- The router LCP is in the Opened state.

Otherwise PPP takes no action on the VSA. If you do not enable the **lcp-connection-update** option, PPP processes the notification from authd, but takes no action.

Default: Disabled

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

Dynamic Profiles Overview

Configuring a PPPoE Dynamic Profile

Configuring Dynamic Authentication for PPP Subscribers | **131**

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces | **126**

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

How to Configure RADIUS-Sourced Connection Status Updates to CPE Devices | **125**

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the PPP Attributes for a Group Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

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 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 | 64

preference (Tunnel Profile)

Syntax

```
preference number;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

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.

This value can be overridden by RADIUS attribute Tunnel-Preference [83].

Options

number—Number that indicates the order in which the router attempts to connect to the destination. Zero is the highest level of preference.

Range: 0 through 2000

Default: 2000

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 | 216

primary-interface (Aggregated Inline Services)

Syntax

```
primary-interface interface-name;
```

Hierarchy Level

[edit interfaces asix aggregated-inline-services-options]

Release Information
Statement introduced in Junos OS Release 16.2.

Description
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 **[edit chassis fpc *slot* pic *number* inline-services bandwidth]** hierarchy level must be the same for both member links.

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.

Options
interface-name—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

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278
Configuring an L2TP LNS with Inline Service Interfaces | 261

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 | jsrc | 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-access-line-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;
            }
        }
        override {
            calling-station-id remote-circuit-id;
            nas-ip-address tunnel-client-gateway-address;
            nas-port tunnel-client-nas-port;
            nas-port-type tunnel-client-nas-port-type;
        }
        remote-circuit-id-delimiter;
        remote-circuit-id-fallback {
        remote-circuit-id-format;
            agent-circuit-id;
            agent-remote-id;
        }
        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-limit-per-username number;
    session-options {
        client-idle-timeout minutes;
        client-idle-timeout-ingress-only;
        client-session-timeoutminutes;
        pcc-context {
            input-service-filter-name filter-name;
            input-service-set-name service-set-name;
            ipv6-input-service-filter-name filter-name;
            ipv6-input-service-set-name service-set-name;
            ipv6-output-service-filter-name filter-name;
            ipv6-output-service-set-name service-set-name;
            output-service-filter-name filter-name;
            output-service-set-name service-set-name;
            profile-name pcef-profile-name;
        }
        strip-user-name {
            delimiter [ delimiter ];
            parse-direction (left-to-right | right-to-left);
        }
    }
    subscriber username {
        delegated-pool delegated-pool-name;
        framed-ip-address ipv4-address;
        framed-ipv6-pool ipv6-pool-name;
        framed-pool ipv4-pool-name;
        password password;
        target-logical-system logical-system-name <target-routing-instance (default | routing-instance-name>;
        target-routing-instance (default | routing-instance-name);
    }
}
```

Hierarchy Level

```
[edit access]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a subscriber access profile that includes 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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.
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 | **268**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring Service Accounting with JSRC*
- *Configuring Service Accounting in Local Flat Files*
- *AAA Service Framework Overview*
- *Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management*

proxy-mode

Syntax

```
proxy-mode;
```

Hierarchy Level

[edit forwarding-options dhcp-relay **overrides**],
[edit forwarding-options dhcp-relay group *group-name* **overrides**],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay **overrides**],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay **overrides**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* **overrides**],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* **overrides**]

Release Information

Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.

You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *DHCP Relay Proxy Overview*
- *Extended DHCP Relay Agent Overview*
- *Enabling DHCP Relay Proxy Mode*

ps0 (Pseudowire Subscriber Interfaces)

Syntax

```
ps0 {
    anchor-point lt-device;
    mtu bytes;
    mac mac-address;
    no-gratuitous-arp-request;
    (flexible-vlan-tagging | stacked-vlan-tagging | untagged | vlan-tagging);
}
```

Hierarchy Level

```
[edit logical-systems transport-ls interfaces]
```

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure the pseudowire logical device.

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

- Pseudowire Subscriber Logical Interfaces Overview | 331
- Configuring a Pseudowire Subscriber Logical Interface | 338
- Configuring a Pseudowire Subscriber Logical Interface Device | 341
- Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface | 346
- Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface | 350

pseudowire-service (Pseudowire Subscriber Interfaces)

Syntax

```
pseudowire-service {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis]
```

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure properties for the pseudowire devices on the router.

The remaining statement is 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

- Pseudowire Subscriber Logical Interfaces Overview | 331
- Configuring a Pseudowire Subscriber Logical Interface | 338
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router | 340

qualified-next-hop (Dynamic Access-Internal Routes)

Syntax

```
qualified-next-hop interface-name {
    mac-address address;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options access-internal route
    subscriber-ip-address],
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name
    access-internal route subscriber-ip-address],
[edit dynamic-profiles profile-name routing-options access-internal route subscriber-ip-address]
```

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 route *subscriber-ip-address*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address*]** 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

interface-name—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. 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 Access-Internal Routes for DHCP and PPP Subscribers | 66

radius (Access Profile)

Syntax

```
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;
            nas-identifier;
        }
        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;
}
ip-address-change-notify message;
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-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;
            }
        }
        override {
            calling-station-id remote-circuit-id;
            nas-ip-address tunnel-client-gateway-address;
            nas-port tunnel-client-nas-port;
            nas-port-type tunnel-client-nas-port-type;
        }
        remote-circuit-id-delimiter;
        remote-circuit-id-fallback;
        remote-circuit-id-format {
            agent-circuit-id;
            agent-remote-id;
        }
        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;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

Options

accounting-server—(MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

Values: *ip-address*—IP version 4 (IPv4) address.

authentication-server—(SRX Series only) Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

Values: *ip-address*—IPv4 address.

preauthentication-server—(MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

NOTE: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the **exclude** statement.

Values: *ip-address*—IPv4 address.

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

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *RADIUS Logical Line Identifier (LLID) Overview*
- *Configuring Logical Line Identification (LLID) Preauthentication*

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.

NOTE: IPv6CP negotiation messages are not rejected for static interfaces.

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

- *Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address*

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. 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

Using DHCP Relay Agent Option 82 Information

dhcp-relay | **489**

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 | 216

report-ingress-shaping-rate (Dynamic CoS Interfaces)

Syntax

```
report-ingress-shaping-rate bps;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*]

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.

> **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.

Options

bps—Ingress shaping rate in bits per second.

Range: 1000 through 6,400,000,000,000 (6.4Tbps)

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 | **244**
- Transmission of Tx and Rx Connection Speeds from LAC to LNS | **236**
- *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

```
request services l2tp destination unlock destination-name
```

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

destination-name—Name of the destination to be unlocked.

Required Privilege Level

view

RELATED DOCUMENTATION

- Removing an L2TP Destination from the Destination Lockout List | **181**
- Configuring the L2TP Destination Lockout Timeout | **180**
- show services l2tp destination lockout | **936**

List of Sample Output

request services l2tp destination unlock on page 744

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

```
retransmission-count-established count;
```

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.

> **BEST PRACTICE:** Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp tunnel retransmission-count-established**.

Options

count—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

- Configuring Retransmission Attributes for L2TP Control Messages | **160**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**

retransmission-count-not-established (L2TP)

Syntax

```
retransmission-count-not-established count;
```

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.

> **BEST PRACTICE:** Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp tunnel retransmission-count-not-established**.

Options

count—Number or retransmissions.

Range: 2 through 30

Default: 5

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Configuring Retransmission Attributes for L2TP Control Messages | **160**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**

route (Access)

Syntax

```
route prefix {
    metric route-cost;
    next-hop next-hop;
    preference route-distance;
    tag route-tag;
    tag2 route-tag2;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access],
[edit dynamic-profiles *profile-name* routing-options access],
[edit dynamic-profiles *profile-name* routing-options rib *routing-table-name* access]

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 access]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the parameters for access routes.

Options

prefix—Either the specific route prefix that you want to assign to the access route or one of the following route prefix variables.

- 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]. When you use this variable, you must configure it at either of the following hierarchy levels:
  - **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib $junos-ipv6-rib access]**
  - **[edit dynamic-profiles *profile-name* routing-options rib $junos-ipv6-rib access]**

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 Access Routes for Subscriber Management | 64

route (Access Internal)

Syntax

```
route subscriber-ip-address {
    next-hop next-hop;
    qualified-next-hop underlying-interface {
        mac-address address;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options access-internal],
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name
    access-internal],
[edit dynamic-profiles profile-name routing-options access-internal]
```

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 access-internal]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure parameters for an access-internal route.

Options

subscriber-ip-address—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.

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

route-suppression (DHCP Local Server and Relay Agent)

Syntax

```
route-suppression (access | access-internal | destination);
```

Hierarchy Level

[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 *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*...],
[edit routing-instances *routing-instance-name* ...],
[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*]

Release Information

Statement introduced in Junos OS Release 13.2.

Description

Configure the jdhcpd process to suppress the installation of access, access-internal, or destination routes during client binding.

> **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.

Options

access—(DHCPv6 only) Suppress installation of access routes. You can use the **access** and **access-internal** options in the same statement for DHCPv6.

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.

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.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

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

routing-instance (Tunnel Profile)

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Specify a routing instance for a tunnel.

Options
routing-instance-name—Name of the routing instance.
Default: Routing instance *default*

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 | 216

routing-instance (L2TP Destination)

Syntax

```
routing-instance routing-instance-name {
    drain;
}
```

Hierarchy Level

```
[edit services l2tp destination address ip-address]
```

Release Information

Statement introduced in Junos OS Release 13.2.

Description

Specify the routing instance in which the destination is created.

Options

routing-instance-name— Name of the routing instance.

Default: Routing instance *default*

The remaining statement is 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

Configuring L2TP Drain | 181

routing-instance (L2TP Tunnel Destination)

Syntax

```
routing-instance routing-instance-name {
    drain;
}
```

Hierarchy Level

```
[edit services l2tp tunnel name tunnel-name address ip-address]
```

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

routing-instance-name— Name of the routing instance.

Default: Routing instance *default*

The remaining statement is 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

Configuring L2TP Drain | 181

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;
                tag2 route-tag2;
            }
        }
        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;
                    tag2 route-tag2;
                }
            }
            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 in a dynamic client profile or a dynamic service profile.

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. 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 a Dynamic Profile for use by a Retailer in the DHCPv4 Solution*

routing-options (Dynamic Profiles)

Syntax

```
routing-options {
    access {
        route prefix {
            metric route-cost;
            next-hop next-hop;
            preference route-distance;
            tag route-tag;
            tag2 route-tag2;
        }
    }
    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;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
    }
```

```
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name],
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance]
```

Release Information

Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]** hierarchy level introduced in Junos OS Release 10.1.

Description

Configure protocol-independent routing properties in a dynamic client 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 Access Routes for Subscriber Management | 64

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers | 66

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.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

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.

NOTE: To create more than one IP reassembly rule, include the **rule** statement multiple times.

Options

rule-name—Name of the IP reassembly rule.

Range: Up to 63 characters

The remaining statement is 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

Configuring IP Inline Reassembly for L2TP | 318

IP Packet Fragment Reassembly for L2TP Overview | 317

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

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS | 244

Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal | 244

rx-window-size (L2TP)

Syntax

```
rx-window-size packets;
```

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

packets—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.

Range: 4 through 128

Default: 4

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Setting the L2TP Receive Window Size | **177**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**

secondary-interface (Aggregated Inline Services)

Syntax

```
secondary-interface interface-name;
```

Hierarchy Level

[edit interfaces asix aggregated-inline-services-options]

Release Information
Statement introduced in Junos OS Release 16.2.

Description

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.

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.

Options

interface-name—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

Configuring 1:1 LNS Stateful Redundancy on Aggregated Inline Service Interfaces | 278

Configuring an L2TP LNS with Inline Service Interfaces | 261

secret (Tunnel Profile)

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Specify the tunnel password sent to the LNS for authentication.

Options
password—Cleartext password.

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 | 216

service-device-pool (L2TP)

Syntax

```
service-device-pool pool-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

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

pool-name—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

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301

service-device-pools (Service Interfaces)

Syntax

```
service-device-pools {
    pool pool-name {
        interface service-interface-name;
    }
}
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement introduced in Junos OS Release 11.4.
Support for **ps** service interfaces added in Junos OS Release 19.3R1.

Description

Configure one or more pools of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The interfaces in the pool must be inline service interfaces (si). The service device pool is required for dynamic LNS sessions.

Options

pool-name—Name of the service interface pool.

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

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions | 313

service-interface (L2TP Processing)

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Release Information

Statement introduced before Junos OS Release 7.4.
si-*fpc*/*pic*/*port* option added in Junos OS Release 11.4.
Option **asi*fpc* added in Junos OS Release 16.2.**

Description

Specify the service interface responsible for handling L2TP processing.

> **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.

Options

interface-name—Name of the service interface. The ae, si, and sp interface types are supported as follows:

- **asix**—(MPCs on MX Series routers) Aggregated inline services interface.
- **sp-*fpc*/*pic*/*port***—On AS or Multiservices PICs on M7i, M10i, and M120 routers.
- **si-*fpc*/*pic*/*port***—On MPCs on MX Series routers.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Local Gateway Address and PIC

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301

service-profile (L2TP)

Syntax

```
service-profile profile-name(parameter)&profile-name;
```

Hierarchy Level

```
[edit access profile profile-name client client-name l2tp]
[edit services l2tp tunnel-group group-name]
```

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers.

Description

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.

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.

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.

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.

Options

profile-name—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 (&).

parameter—(Optional) Value to be passed to the service when it is activated on the subscriber session.

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 Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**

Configuring an L2TP Access Profile on the LNS | **268**

service-rate-limiter (Access)

Syntax

```
service-rate-limiter {
    rx-multiplier rx-multiplier;
    service-name service-profile-name;
    tx-multiplier tx-multiplier;
}
```

Hierarchy Level

```
[edit access]
```

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers.

Description

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).

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.

Options

rx-multiplier—(Optional) Multiplier applied to convert the Rx connect speed value Kbps to bps and optionally adjust the rate up or down.

Default: 1000

Range: 1 through 2000

service-profile-name—Name of the dynamic service profile conveyed in VSA 26-65 that specifies upstream and downstream traffic rates.

tx-multiplier—(Optional) Multiplier applied to convert the Tx connect speed value Kbps to bps and optionally adjust the rate up or down.

Default: 1000

Range: 1 through 2000

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds | **255**
- Configuring an L2TP LAC | **183**
- Transmission of Tx and Rx Connection Speeds from LAC to LNS | **236**

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 *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 session mode.

Options

Default: automatic

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.

multihop—Configure multihop BFD sessions and passive DHCP clients.

single-hop—Configure single hop BFD sessions and non-passive DHCP clients.

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

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

session-options

Syntax

```
session-options {
    client-group [ group-names ];
    client-idle-timeout minutes;
    client-idle-timeout-ingress-only;
    client-session-timeoutminutes;
    pcc-context {
        input-service-filter-name filter-name;
        input-service-set-name service-set-name;
        ipv6-input-service-filter-name filter-name;
        ipv6-input-service-set-name service-set-name;
        ipv6-output-service-filter-name filter-name;
        ipv6-output-service-set-name service-set-name;
        output-service-filter-name filter-name;
        output-service-set-name service-set-name;
        profile-name pcef-profile-name;
    }
    strip-user-name {
        delimiter [ delimiter ];
        parse-direction (left-to-right | right-to-left);
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 8.5.

Description

(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.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

(MX Series) Specify characteristics related to policy and charging control (PCC) rules, such as the PCEF profile that contains the rules, service sets to process the rules, and service filters for the service sets.

Options

client-idle-timeout—(MX Series only) Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

During this period, the router determines whether the subscriber is inactive by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, non-DHCP subscribers (such as L2TP or PPP) are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout; DHCP subscribers are disconnected.

When you additionally configure the related **client-idle-timeout-ingress-only** statement (MX Series only), the router monitors only ingress traffic to determine whether the subscriber is inactive; it does not monitor any egress traffic. The related **client-session-timeout** statement terminates the subscriber session when the session timeout expires regardless of user activity.

Client idle timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model. It is not practical for DHCP or DHCPv6 subscribers.

Although you can use the **client-idle-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the idle timeout for VLANs, the timeout period starts when the VLAN is instantiated. It resets when a client session is created or an existing session is reactivated. When no traffic is detected on an authenticated VLAN for the duration of the timeout, the VLAN is considered inactive and is deleted. If no client sessions are ever created on the VLAN, then the VLAN is removed when the timeout expires.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes of idle time that elapse before the session is terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 10 through 1440 minutes

client-idle-timeout-ingress-only—(MX Series only) Starting in Junos OS Release 16.2, specify that only ingress traffic is monitored for subscriber idle timeout processing for the duration of the idle timeout period that you specify with the **client-idle-timeout** statement. If no ingress traffic is received for the duration of the timeout, then the subscriber is gracefully logged out (non-DHCP subscribers) or disconnected (DHCP subscribers).

If you configure **client-idle-timeout** alone, then both ingress and egress traffic are monitored during the idle timeout. Monitoring only ingress traffic is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore does not detect that the peer is not up. Because the LAC monitors both ingress and egress traffic by default, in this situation it receives the egress traffic from the LNS and either

does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored in this case, the LAC can detect that the peer is inactive and then initiate logout.

client-session-timeout—(SRX Series, vSRX only) Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Alternatively, when you want subscribers to be identified as inactive before they are terminated, use the related statements, **client-idle-timeout** and **client-idle-timeout-ingress-only**. Use **client-idle-timeout** alone to specify a period of time during which both ingress and egress subscriber data traffic is monitored; if no traffic is detected for the duration of the period, the subscriber is considered inactive and is terminated. Add the **client-idle-timeout-ingress-only** statement to monitor only ingress traffic for the duration of the timeout set with the **client-idle-timeout statement**.

BEST PRACTICE: We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is a simple time-based timeout, it is likely to interrupt subscribers actively using a voice service and terminate their calls unexpectedly (from the subscriber viewpoint). This result is a particular concern for emergency services calls.

Client session timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model when no voice services are offered. For DHCP or DHCPv6 subscribers, the session timeout is used as the DHCP lease timer if no other lease time configuration is present.

Although you can use the **client-session-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the session timeout for VLANs, the timeout period starts when the VLAN is instantiated.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes after which user sessions are terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 1 through 527040 minutes

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

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Subscriber Session Timeout Options*
- *Configuring Username Modification for Subscriber Sessions*
- *Removing Inactive Dynamic Subscriber VLANs*
- *Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management*

sessions-limit-group (L2TP)

Syntax

```
sessions-limit-group limit-group-name {
    maximum-sessions number;
}
```

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

limit-group-name—Identifier of the session-limit group for which the session limit is configured.

The remaining statement is 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

- Limiting the Number of L2TP Sessions Allowed by the LAC or LNS | **213**
- Configuring an L2TP LAC | **183**
- Configuring an L2TP LNS with Inline Service Interfaces | **261**
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | **301**
- L2TP for Subscriber Access Overview | **153**

soft-gre

Syntax

```
soft-gre group-name {
    destination-networks [prefix] {
    dynamic-profile profile-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.

dynamic-profile *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 ps*n*—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 | **360**
- Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway | **361**
- Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation | **365**
- show services soft-gre tunnel | **988**

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 | 216

stacked-vlan-tagging

Syntax

```
stacked-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro 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

- *Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*

statistics (Access Profile)

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name 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

Configuring Authentication and Accounting Parameters for Subscriber Access

strip-user-name (Access Profile)

Syntax

```
strip-user-name {
    delimiter delimiter;
    parse-direction (left-to-right | right-to-left);
}
```

Hierarchy Level

```
[edit access profile profile-name session-options]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

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.

You can use the **aaa-options** ***aaa-options-name*** 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.

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

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Username Modification for Subscriber Sessions*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | **266**
- Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | **263**

subscriber-context (AAA Options)

Syntax

```
subscriber-context subscriber-context-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

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.

NOTE: Only the default (master) logical system is supported.

Options

subscriber-context-name—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

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile | 266

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface | 263

subscriber-management (Subscriber Management)

Syntax

```
subscriber-management {
    enable;
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    }
    overrides {
        event {
            catastrophic-failure {
                reboot (master | standby);
            }
        }
        interfaces {
            family (inet | inet6) {
                layer2-liveness-detection;
                ipoe-dynamic-arp-enable;
                receive-gratuitous-arp;
            }
        }
        no-unsolicited-ra;
        ra-initial-interval-max seconds;
        ra-initial-interval-min seconds;
        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);
        |
    }
    redundancy {
        interface name {
            local-inet-address v4-address;
            local-inet6-address v6-address;
            shared-key string;
```

```
        virtual-inet-address virtual-v4-address;
        virtual-inet6-address virtual-v6-address;
      }
      no-advertise-routes-on-backup;
      protocol {
        pseudo-wire;
        vrrp;
      }
    }
    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 Maintain DHCP Subscribers During Interface Delete Events

Tracing Subscriber Management Database Events for Troubleshooting

Configuring Junos OS Enhanced Subscriber Management

Configuring Junos OS Enhanced Subscriber Management

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover | **62**

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization

tag (Access)

Syntax

```
tag route-tag;
```

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 10.2.
Support at the **[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 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:

- **$junos-framed-route-tag**—Tag assigned to an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-Route attribute [22].
- **$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 | 64

tag2 (Dynamic Access Routes)

Syntax

```
tag2 route-tag2;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options **access route** *prefix*],
[edit dynamic-profiles *profile-name* routing-options **access route** *prefix*]

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Options

route-tag2—One of the following values the specific tag2 value you want to assign to the access route or the following predefined variable:

- A specific tag 2 value for the specified access route prefix.
- **$junos-framed-route-tag2**—Tag2 value assigned to an IPv4 access route. The value is dynamically replaced with the preference value (subattribute 6) from the RADIUS Framed-Route attribute [22]. You configure this variable only when the access route prefix is derived from the $junos-framed-route-ip-address-prefix predefined variable; this value is (subattribute 1) of the RADIUS Framed-Route attribute [22].

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Migrating Static PPP Subscriber Configurations to Dynamic Profiles Overview | 127

threshold (detection-time)

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[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 group-name liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd detection-time]
```

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

> **NOTE:** The threshold time must be greater than or equal to the **minimum-interval** or the **minimum-receive-interval.**

Options

milliseconds— Value for the detection time adaptation threshold.

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

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

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

threshold (transmit-interval)

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[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 group-name liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd transmit-interval]
```

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

milliseconds — Threshold value.

Range: 0 through 4,294,967,295 (2^{32} – 1)

> **NOTE:** The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

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

tos-reflect (L2TP)

Syntax

```
tos-reflect;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

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 *interface-name* unit *logical-unit-number* 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

- *Configuring Dynamic CoS for an L2TP LNS Inline Service*

trace (DHCP Relay Agent)

Syntax

```
trace;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name interface
    interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
    dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
    group group-name interface interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-name interface
    interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name interface
    interface-name]
```

Release Information

Statement introduced in Junos OS Release 10.4.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

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.

EX Series switches do not support DHCPv6.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Configuring DHCP Relay Agent*
- *Tracing Extended DHCP Operations*

Tracing Extended DHCP Operations for Specific Interfaces

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: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Tracing L2TP Operations

Tracing L2TP Events for Troubleshooting | 323

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-expresion;
        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: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

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 PPP Service Events for Troubleshooting | 145

traceoptions (Subscriber Management)

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag;
}
```

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

file *filename*—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 **/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

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.
- **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.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

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: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

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 Events for Troubleshooting*

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. 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

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

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

tunnel (L2TP)

Syntax

```
tunnel {
    assignment-id-format (assignment-id | client-server-id);
    failover-resync (failover-protocol | silent-failover);
    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. 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

Configuring an L2TP LAC | **183**

Configuring an L2TP LNS with Inline Service Interfaces | **261**

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 | 216

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. 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

Configuring L2TP Tunnel Groups

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces | 301

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 | 175

tunnel-profile (Tunnel Profile)

Syntax

```
tunnel-profile profile-name {
    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]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Define a tunnel profile for subscriber access.

Options

profile-name—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].

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

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 | 216

tunnel-switch-profile (L2TP Tunnel Switching, Application)

Syntax

```
tunnel-switch-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit services l2tp],
[edit services l2tp tunnel-group group-name]
```

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
profile-name—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

Configuring L2TP Tunnel Switching | 175

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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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 | 175

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
Configure whether the LAC accepts, ignores, or rejects requests from an LNS to change the destination IP address, UDP port, or both.

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.

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.

Default
The LAC accepts IP address or UDP port change requests from the LNS.

Options
accept—Accept all change requests for IP address or UDP port.

ignore—Ignore all change requests for IP address or UDP port.

ignore-ip-address—Ignore a change request for IP address, but accept a change request for UDP port.

ignore-udp-port—Ignore a change request for UDP port, but accept a change request for IP address.

reject—Reject all change requests for IP address or UDP port.

reject-ip-address—Reject a change request for IP address, but accept a change request for UDP port.

reject-udp-port—Reject a change request for UDP port, but accept a change request for IP address.

Required Privilege Level
interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring How the LAC Responds to Address and Port Changes Requested by the LNS | 186

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 236 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 | **244**

Transmission of Tx and Rx Connection Speeds from LAC to LNS | **236**

type (Tunnel Profile)

Syntax

```
type tunnel-type;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Specify the tunnel type (protocol).

Default
l2tp

Options
tunnel-type—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

Configuring a Tunnel Profile for Subscriber Access | 216

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** *profile-name* **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

logical-unit-number—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the **unit *logical-unit-number*** statement for dynamic PPPoE logical interfaces, you must use the predefined variable **$junos-interface-unit** in place of ***logical-unit-number***. The **$junos-interface-unit** predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.

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

Configuring a PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

unit (Dynamic Profiles Standard Interface)

Syntax

```
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;
        }
    }
    dial-options {
        ipsec-interface-id name;
        l2tp-interface-id name;
        (shared | dedicated);
    }
    encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid
        | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
        ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp
        | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc | 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);
    family family {
        address address;
        demux-destination,
        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 {
        fail-filter filter-name;
        mode loose;
    }
    service {
        input {
            service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        input-vlan-map {
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (push | swap);
            tag-protocol-id tpid;
            vlan-id number;
        }
        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
        output-vlan-map {
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (pop | swap);
            tag-protocol-id tpid;
            vlan-id number;
        }
    }
    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 {
        shared-name  filter-shared-name;
    }
    output filter-name {
        shared-name  filter-shared-name;
    }
}
host-prefix-only;
keepalives {
    interval seconds;
    }
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
service {
    pcef  pcef-profile-name {
        activate rule-name | activate-all;
    }
}
```

```
    targeted-options {
        backup backup;
        group group;
        primary primary;
        weight ($junos-interface-target-weight | weight-value);
    }
    vlan-id number;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following predefined variables:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

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

- *Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Agent Circuit Identifier-Based Dynamic VLANs 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
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Pseudowire Subscriber Logical Interface Device | 341

username-include (Local Authentication)

Syntax

```
username-include {
    circuit-id;
    delimiter character;
    domain-name name;
    mac-address;
    remote-id;
}
```

Hierarchy Level

[edit dynamic-profiles *name* interfaces $junos-interface-ifd-name **unit** $junos-interface-unit **ppp-options local-authentication**]

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Configure a local username that authd can use to request authentication for terminated PPP subscribers. This enables the external RADIUS server to pass implementation-specific configuration for successfully authenticated subscribers. Local authentication supports CPEs that do not negotiate authentication protocols in the same dynamic profile as CPEs that use only PAP or CHAP authentication.

The username takes the following format when you use the default delimiter:

mac-address.circuit-id.remote-id@domain-name

Options

circuit-id—Include the agent circuit identifier (ACI) in the local username.

delimiter *character*—Specify the character that separates components that make up the concatenated username.

Default: Period (.)

domain-name *name*—Specify the domain name that ends the local username created for the subscribers. The username is sent to RADIUS in the Access-Request message. The string can include the following characters: a through z, A through Z, 0 through 9, “-“, or “.”.

mac-address—Include the MAC address from the client PDU in the local username.

remote-id—Include the agent remote identifier (ARI) in the local username.

Required Privilege Level
interface

RELATED DOCUMENTATION

Configuring Local Authentication in Dynamic Profiles for Static Terminated IPv4 PPP Subscribers | **128**

version (BFD)

Syntax

```
version (0 | 1 | automatic);
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp oam bfd-liveness-detection],
[edit logical-systems *logical-system-name* protocols ldp oam fec *address* bfd-liveness-detection],
[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**],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec *address* bfd-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 BFD protocol version to detect.

Options

0—Use BFD protocol version 0.

1—Use BFD protocol version 1.

automatic—Autodetect the BFD protocol version.

Default: automatic

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

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

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
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

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.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Weighted Load Balancing for LAC Tunnel Sessions | **221**

Configuring the L2TP LAC Tunnel Selection Parameters | **219**

vlan-id (Dynamic Profiles)

Syntax

```
vlan-id (number | none);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

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

number—A valid VLAN identifier. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN identifier.

none—Enable the use of untagged pseudo-wire frames on dynamic interfaces.

- 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

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles

vlan-tagging

Syntax

```
vlan-tagging;
```

Syntax (QFX Series, NFX Series, and EX4600)

```
vlan-tagging;
```

Syntax (SRX Series Interfaces)

```
vlan-tagging native-vlan-id vlan-id;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name]
```

QFX Series, NFX Series, and EX4600 Interfaces

```
[edit interfaces (QFX Series) interface-name ]
[edit interfaces (QFX Series) interface-range interface-range-name ]
```

SRX Series Interfaces

```
[edit interfaces interface ]
```

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 9.5.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro 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.

> **NOTE:** For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.
>
> 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.

Default

VLAN tagging is disabled by default.

Options

native-vlan-id— (SRX Series)Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

> **NOTE:** The **native-vlan-id** can be configured only when either **flexible-vlan-tagging** mode or **interface-mode** trunk is configured.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *802.1Q VLANs Overview*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Tagged Aggregated Ethernet Interfaces*
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *vlan-id*
- *Configuring a Layer 3 Logical Interface*
- *Configuring VLAN Tagging*

vlan-tagging (Dynamic)

Syntax

```
vlan-tagging;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* **interfaces** *interface-name*],
[edit **interfaces** *interface-name*]

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.

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.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

Configuring the L2TP LNS Peer Interface | 273

vlan-tags

Syntax

```
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

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 **[edit interfaces *interface-name*]** hierarchy level.

> **NOTE:** The **inner-range *vid1–vid2*** option is supported on IQE PICs only.

Options

inner [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN ID.

> **NOTE:** On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range *tpid. vid1—vid2*** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-stacked-vlan-id** 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

Configuring Dual VLAN Tags

CHAPTER 9

Operational Commands

clear services l2tp destination | **844**

clear services l2tp destination lockout | **846**

clear services l2tp session | **848**

clear services l2tp session statistics | **851**

clear services l2tp tunnel | **853**

clear services l2tp tunnel statistics | **856**

request interface (revert | switchover) (Aggregated Inline Service Interfaces) | **858**

show ancp subscriber | **860**

show bfd subscriber session | **870**

show dynamic-profile session | **876**

show interfaces ps0 (Pseudowire Subscriber Interfaces) | **882**

show interfaces redundancy | **888**

show ppp interface | **892**

show ppp statistics | **912**

show ppp summary | **921**

show services inline ip-reassembly statistics | 923

show services l2tp client | 929

show services l2tp destination | 931

show services l2tp destination lockout | 936

show services l2tp session | 938

show services l2tp session-limit-group | 949

show services l2tp summary | 951

show services l2tp tunnel | 959

show services l2tp tunnel-group | 967

show services l2tp tunnel-switch destination | 969

show services l2tp tunnel-switch session | 974

show services l2tp tunnel-switch summary | 980

show services l2tp tunnel-switch tunnel | 982

show services soft-gre tunnel | 988

show subscribers | 992

show subscribers summary | 1040

show system subscriber-management statistics | 1049

show system subscriber-management summary | 1059

test services l2tp tunnel | 1063

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 | 931

List of Sample Output
clear services l2tp destination all on page 845

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 lockout

Syntax

```
clear services l2tp destination lockout
<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 lockout 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 lockout timer removes the destination from the lockout 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 | 844

show services l2tp destination | 931

List of Sample Output

clear services l2tp destination lockout all on page 847

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 | **851**

show services l2tp session | **938**

List of Sample Output
clear services l2tp session on page 850
clear services l2tp session interface on page 850

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  Remote   State               Interface         Interface
  ID     ID                           unit              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  Remote   State               Interface         Interface
   ID     ID                           unit              Name
   34915  2        Established         1073741829        si-2/1/0
   46142  4        Established         1073741831        si-2/1/0
```

clear services l2tp session statistics

Syntax

clear services l2tp session statistics (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.
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

all—Clear statistics for all L2TP sessions.

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 statistics for only the L2TP sessions associated with the local gateway with the specified address.

local-gateway-name *gateway-name*—Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.

local-session-id *session-id* —Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id *tunnel-id*—Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.

tunnel-group *group-name*—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.

user *username* —Clear statistics for only the L2TP sessions for the specified username. 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 session | **848**
- show services l2tp session | **938**

List of Sample Output
clear services l2tp session statistics all on page 852

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 | **856**

show services l2tp tunnel | **959**

List of Sample Output
clear services l2tp tunnel on page 854

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 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 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

all—Clear statistics for all L2TP tunnels.

interface sp-*fpc/pic/port*—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.

local-gateway *gateway-address*—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*—Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.

tunnel-group *group-name*—Clear statistics for only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

clear

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

clear services l2tp tunnel | **853**

show services l2tp tunnel | **959**

List of Sample Output

clear services l2tp tunnel statistics all on page 857

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 859
request interface revert on page 859

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
```

show ancp subscriber

Syntax

```
show ancp subscriber
<brief | detail>
<access-aggregation-circuit-id circuit-identifier>
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Release Information

Command introduced in Junos OS Release 9.4.
neighbor option replaced with **ip-address** in Junos OS Release 16.1.
system-name option introduced in Junos OS Release 16.1.
access-aggregation-circuit-id option introduced in Junos OS Release 18.4R1.

Description

Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACI, or the subscriber associated with the specified ANCP neighbor (access node).

After an ancpd restart, this command displays orphaned entries (marked with an **o**) for subscriber sessions that were established before the restart but which have not yet been reestablished. As sessions are reestablished, the number of orphaned entries displayed by the command decreases. The number reaches zero when all sessions are reestablished or when the orphaned-interface timer expires.

Options

none—Display information about all subscribers.

brief | detail—(Optional) Display the specified level of detail.

access-aggregation-circuit-id *circuit-identifier*—(Optional) Display information about ANCP subscribers whose Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) matches the specified value.

A *circuit-identifier* that begins with the**#** character indicates a backhaul line identifier. You can specify a wildcard (*) anywhere in the string.

identifier *identifier*—(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.

ip-address *ip-address* —(Optional) Display information about the subscribers connected to the access node specified by the IP address.

system-name *mac-address*—(Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

RELATED DOCUMENTATION

- *clear ancp subscriber*
- *show ancp cos*
- *show ancp neighbor*
- *show ancp statistics*

List of Sample Output

show ancp subscriber on page 865
show ancp subscriber (After ancpd Restart) on page 866
show ancp subscriber brief on page 866
show ancp subscriber detail on page 866
show ancp subscriber access-aggregation-circuit-id detail on page 868
show ancp subscriber identifier identifier-string detail on page 868

Output Fields

Table 33 on page 862 lists the output fields for the **show ancp subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 33: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
Loop Identifier	Access loop identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established. An **o** indicates that the entry is for an orphaned interface and represents a previously established subscriber session that has not been reestablished after an ancpd restart. The number of orphaned entries decreases as the ANCP neighbors reestablish adjacencies and the protocol subscriber sessions are reestablished. The command output indicates this by removing the **o** marker. Eventually the number of orphaned entries reaches zero, because either all the adjacencies and subscriber sessions have been reestablished or any remaining orphaned entries are removed when the orphaned-interface timer expires.	**brief** none
DSL Line State	State of the DSL line: **Idle**, **Showtime**, or **Silent**.	**brief detail**
Access Type	Type of access line employed by the access node: **ADSL1**, **ADSL2**, **ADSL2+**, **VDSL1**, **VDSL2**, **SDSL**, **G.fast**, **VDSL2 Annex Q**, **SDSL bonded**, **VDSL2 bonded**, **G.fast bonded VDSL2 Annex Q bonded** or **OTHER**.	**brief detail** none
Interface	Name of the interface set or logical interface.	**brief detail** none
Rate Kbps	Actual downstream data rate for this local loop.	**brief** none
Neighbor	IP address of ANCP neighbor (access node).	**brief** none

Table 33: show ancp subscriber Output Fields *(continued)*

Field Name	Field Description	Level of Output
Access Loop Circuit Identifier	Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.	**detail**
Neighbor IP Address	IP address of the ANCP neighbor (access node).	**detail**
Aggregate Circuit Identifier	ASCII identifier for the subscriber access loop; value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003).	**detail**
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	**detail**
Tech Type	Type of technology employed by the subscriber. Currently Junos OS supports **DSL** technology type only.	**detail**
DSL Line Data Link	Data link protocol employed on the access loop: **AAL5** or **Ethernet**.	**detail**
DSL Line Encapsulation	Encapsulation type on the access loop, for Ethernet only: • **0**—NA, type not conveyed • **1**—Untagged Ethernet • **2**—Single-tagged Ethernet	**detail**

Table 33: show ancp subscriber Output Fields *(continued)*

Field Name	Field Description	Level of Output
DSL Line Encapsulation Payload	Payload carried across the access loop: • **0**—NA, type not conveyed • **1**—PPPoA LLC • **2**—PPPoA null • **3**—IPoA LLC • **4**—IPoA null • **5**—Ethernet over AAL5 LLC with FCS • **6**—Ethernet over AAL5 LLC without FCS • **7**—Ethernet over AAL5 null with FCS • **8**—Ethernet over AAL5 null without FCS	**detail**
Interface Type	Type of interface employed for subscriber traffic: **ifl** for a single VLAN or **interface-set** for a configured group of VLANs.	**detail**
Actual Net Data Upstream	Actual upstream data rate for this local loop, in Kbps.	**detail**
Actual Net Data Downstream	Actual downstream data rate for this local loop, in Kbps.	**detail**
Minimum Net Data Upstream	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Minimum Net Data Downstream	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Maximum Net Data Upstream	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Maximum Net Data Downstream	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Attainable Net Data Upstream	Maximum attainable upstream data rate for this local loop, in Kbps.	**detail**

Table 33: show ancp subscriber Output Fields *(continued)*

Field Name	Field Description	Level of Output
Attainable Net Data Downstream	Maximum attainable downstream data rate for this local loop, in Kbps.	**detail**
Minimum Low Power Data Downstream	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	**detail**
Minimum Low Power Data Upstream	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	**detail**
Maximum Interleave Delay Downstream	Maximum interleaving delay for downstream data, in milliseconds.	**detail**
Maximum Interleave Delay Upstream	Maximum interleaving delay for upstream data, in milliseconds.	**detail**
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	**detail**
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	**detail**

Sample Output

show ancp subscriber

`user@host>` **show ancp subscriber**

```
  Loop Identifier        DSL Line  Tech Type     Access Type   Interface
Rate   Neighbor
                         State                                 Kbps
**circuit 101            Idle      DSL       ADSL1    ----             32
203.0.113.13
**circuit 102            Idle      DSL       ADSL1    ----             32
```

```
203.0.113.13
  circuit 301            Showtime  DSL       ADSL1    ----              32
203.0.113.15
  circuit 302            Showtime  DSL       ADSL1    ----              32
203.0.113.15
```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
```

```
  Loop Identifier        DSL Line  Tech Type     Access Type  Interface
Rate     Neighbor
                         State                                Kbps
o circuit 201            Showtime  DSL        ADSL1    ----              222222

o circuit 202            Showtime  DSL        ADSL1    ----              222222
```

show ancp subscriber brief

```
user@host> show ancp subscriber brief
```

```
  Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10               VDSL2    set-ge-10410       64       203.0.113.102

  port-1-11               VDSL2    set-ge-10411       64       203.0.113.111

  port-2-10               VDSL2    ge-1/0/4.12        64       203.0.113.112
  port-2-11               VDSL2    ge-1/0/4.13        64       203.0.113.113
```

show ancp subscriber detail

```
user@host> show ancp subscriber detail
```

```
 Subscriber Information
*  Access Loop Circuit Identifier : circuit 101
     Neighbor IP Address                 : 203.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                  : DSL
```

```
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
*  Access Loop Circuit Identifier: circuit 102
     Neighbor IP Address                  : 213.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                  : DSL
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
```

```
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
...
```

show ancp subscriber access-aggregation-circuit-id detail

user@host> **show ancp subscriber access-aggregation-circuit-id "#TEST-DPU-C-100" detail**

```
Subscriber Information
 * Access Loop Circuit Identifier : circuit 201
     Neighbor IP Address                  : 192.0.2.1
     Access Loop Remote Identifier        : remote 123
     Aggregate Circuit Identifier         : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                           : DSL
     Interface Type                       : interface
     Interface                            : ge-1/0/0.3221225475
     Actual Net Data Upstream             : 1024
     Actual Net Data Downstream           : 2048
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
* Access Loop Circuit Identifier : circuit 202
     Neighbor IP Address                  : 192.0.2.1
     Access Loop Remote Identifier        : remote 185
     Aggregate Circuit Identifier         : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                           : DSL
     Interface Type                       : interface
     Interface                            : ge-1/0/0.3221225476
     Actual Net Data Upstream             : 1024
     Actual Net Data Downstream           : 2048
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
```

show ancp subscriber identifier identifier-string detail

user@host> **show ancp subscriber identifier port-1-11 detail**

```
      Access Loop Identifier : port-1-11
          Neighbor IP Address                  : 203.0.113.112
          Aggregate Circuit Identifier Binary  : 0/0
```

```
DSL Type                              :  DSL  0
Interface Type                        : interface-set
Interface                             : set-ge-10411
DSL Line State                        : Show Time
Actual Net Data Upstream              : 64
Actual Net Data Downstream            : 64
DSL Line Data Link                    : AAL5
DSL Line Encapsulation                : N/A
DSL Line Encapsulation Payload        : N/A
Minimum Net Data Upstream             : 64
Minimum Net Data Downstream           : 64
Maximum Net Data Upstream             : 64
Maximum Net Data Downstream           : 64
Attainable Net Data Upstream          : 64
Attainable Net Data Downstream        : 64
Minimum Low Power Data Downstream     : 64
Minimum Low Power Data Upstream       : 64
Maximum Interleave Delay Downstream   : 50
Maximum Interleave Delay Upstream     : 50
Actual Interleave Delay Downstream    : 50
Actual Interleave Delay Upstream      : 50
```

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

> *clear bfd session*
>
> *Examples: Configuring BFD for Static Routes*

List of Sample Output
show bfd subscriber session on page 872
show bfd subscriber session brief on page 872
show bfd subscriber session detail on page 873
show bfd subscriber session extensive on page 873
show bfd subscriber session summary on page 875

Output Fields
Table 34 on page 870 describes the output fields for the **show bfd subscriber session** command. Output fields are listed in the approximate order in which they appear.

Table 34: show bfd subscriber session Output Fields

Field Name	Field Description	Level of Output
Address	IP Address on which the BFD subscriber session is active.	All levels

Table 34: show bfd subscriber session Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the BFD subscriber session: **Up**, **Down**, **Init** (initializing), or **Failing**.	All levels
Interface	Interface on which the BFD subscriber session is active.	All levels
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	All levels
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	All levels
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	All levels
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**.	**detail extensive**
routing table index	Value of the routing table index. A value of 0 (zero) denotes a route in the default routing table managed by enhanced subscriber management.	**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**
RX interval	Receive interval being used because of adaptation.	extensive

Table 34: show bfd subscriber session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	**extensive**
minimum 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**
min RX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	**extensive**
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**

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        demux0.3221225503  90.000  30.000        3
203.0.113.6          Up        demux0.3221225504  90.000  30.000        3
203.0.113.10         Up        demux0.3221225505  90.000  30.000        3
203.0.113.14         Up        demux0.3221225506  90.000  30.000        3
203.0.113.18         Up        demux0.3221225507  90.000  30.000        3
```

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
```

```
                                              Detect  Transmit
Address              State   Interface        Time    Interval  Multiplier
203.0.113.2          Up      demux0.3221225503  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.6          Up      demux0.3221225504  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.10         Up      demux0.3221225505  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.14         Up      demux0.3221225506  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.18         Up      demux0.3221225507  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
```

show bfd subscriber session extensive

```
user@host> show bfd subscriber session extensive
```

```
                                              Detect   Transmit
Address                 State     Interface           Time     Interval
Multiplier
203.0.113.2             Up        demux0.3221225503    90.000    30.000        3

 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 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 28, remote discriminator 1073741825
```

```
                                                    Detect   Transmit
Address                 State     Interface         Time     Interval
Multiplier
203.0.113.6             Up        demux0.3221225504    90.000    30.000        3

 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 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 29, remote discriminator 1073741826

                                                    Detect   Transmit
Address                 State     Interface         Time     Interval
Multiplier
203.0.113.10            Up        demux0.3221225505    90.000    30.000        3

 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 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 30, remote discriminator 1073741827

                                                    Detect   Transmit
Address                 State     Interface         Time     Interval
Multiplier
203.0.113.14            Up        demux0.3221225506    90.000    30.000        3

 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 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 31, remote discriminator 1073741828

                                                    Detect   Transmit
Address                 State     Interface         Time     Interval
```

```
Multiplier
203.0.113.18             Up        demux0.3221225507    90.000    30.000        3

 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 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 32, remote discriminator 1073741829
```

show bfd subscriber session summary

user@host> **show bfd subscriber session summary**

```
                                                Detect  Transmit
Address              State     Interface        Time    Interval  Multiplier
203.0.113.2          Up        demux0.3221225503  90.000  30.000        3
203.0.113.6          Up        demux0.3221225504  90.000  30.000        3
203.0.113.10         Up        demux0.3221225505  90.000  30.000        3
203.0.113.14         Up        demux0.3221225506  90.000  30.000        3
203.0.113.18         Up        demux0.3221225507  90.000  30.000        3
```

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 877

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile) on page 879
show dynamic-profile session service-id (Service Session) on page 880

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

Pseudowire Subscriber Logical Interfaces Overview | 331

List of Sample Output

show interfaces ps0 on page 885
show interfaces ps0 extensive on page 886

Output Fields

Table 35 on page 882 lists the output fields for the **show interfaces ps0** command. Output fields are listed in the approximate order in which they appear.

Table 35: 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 *Common Output Fields Description*.	**brief detail extensive** none

Table 35: show interfaces ps0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	**brief detail extensive** none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	**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: *year-month-day hours*:*minutes*:*seconds*: *timezone* (*hours*:*minutes*:*seconds* ago). 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

Table 35: show interfaces ps0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	**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**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. **NOTE:** The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **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	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. **NOTE:** The packet and byte counts in these fields include traffic that is dropped and does not leave the router.	**detail extensive**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none

Table 35: show interfaces ps0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**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 *Common Output Fields Description*.	**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: Sendbcast-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: Sendbcast-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 889
show interfaces redundancy (Aggregated Ethernet) on page 890
show interfaces redundancy (Aggregated Inline Service Interface) on page 890
show interfaces redundancy detail on page 890

Output Fields

Table 36 on page 889 lists the output fields for the **show interfaces redundancy** command. Output fields are listed in the approximate order in which they appear.

Table 36: 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
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: • 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                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
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

```
show ppp interface interface-name
<extensive | terse>
```

Release Information
Command introduced in Junos OS Release 7.5.

Description
Display information about PPP interfaces.

Options
interface-name—Name of a logical interface.

> Starting in Junos OS Release 17.3, the * (asterisk) wildcard character is supported for the interface name for debugging purpose. With this support, you can match any string of characters in that position in the interface name. For example, so* matches all SONET/SDH interfaces.

extensive | terse—(Optional) Display the specified level of output.

Required Privilege Level
view

List of Sample Output
show ppp interface on page 909
show ppp interface extensive (LCP Connection Update Negotiation Successful) on page 909
show ppp interface extensive (LCP Connection Update Negotiation Failed) on page 910
show ppp interface extensive (Inline Service Interface) on page 910
show ppp interface terse on page 911

Output Fields
Table 37 on page 892 lists the output fields for the **show ppp interface** command. Output fields are listed in the approximate order in which they appear.

Table 37: 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

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
Keepalive settings	Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS-based PPP sessions are supported only on service interfaces (si). • **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	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. • 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**

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
RE Keepalive statistics	Keepalive statistics for the packets handled by the Routing Engine. • **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 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
LCP		**extensive**

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	LCP information:
	• **State**—LCP protocol state (all platforms except M120 and M320 routers):
	  ◦ **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):
	  ◦ **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.
	• **Last updated**—Reports the timestamp of the last successful connection update exchange.
	  1. When LCP negotiation completes, this field has the same value as the **Last completed** field.
	  2. The field then reports the timestamp of any subsequent successful exchange of Connection-Update-Request and Connection-Update-Ack messages with	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
	the peer (such as a home gateway). This field is displayed only when the Connection-Status-Message option is successfully negotiated.	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	• **Negotiated options**:
	  • **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.
	  • **Connection Update Requests**—Number of connection update requests sent by PPP to the remote peer (such as a home gateway). This value does not include retries.
	This field is displayed even when negotiation fails for the Connection-Status-Message option. This enables you to confirm that an update request was sent. The absence of the Juniper Connection Status Message field indicates the peer does not support the updates.
	  • **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.
	  • **Juniper Connection Status Message**—The content of the Connection-Status-Message VSA (26-4874–218) most recently received from RADIUS.
	This field is displayed only when the Connection-Status-Message option is successfully negotiated.
	  • **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	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
	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 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Authentication	CHAP or PAP authentication state information. For CHAP authentication: • **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. For PAP authentication: • **Pap-resp-sent**—PAP response sent to peer (ACK/NACK)t. • **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 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPCP		**extensive**

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	Internet Protocol Control Protocol (IPCP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  ◦ **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**—(M120 and M320 routers) One of the following values:
	  ◦ **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**—IPCP state start time.
	• **Last completed**—IPCP state authentication completion time.
	• **Negotiated options:**
	  ◦ **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.	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
	• **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**	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPV6CP		**extensive**

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	Internet Protocol version 6 Control Protocol (IPv6CP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  ◦ **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**—(M120 and M320 routers) One of the following values:
	  ◦ **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**—IPV6CP state start time.
	• **Last completed**—IPV6CP state authentication completion time.
	• **Negotiated options**:
	  ◦ **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	

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
	quad notation. • **Negotiation mode**—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: **Active** or **Passive**	

OSINLCP State	OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):
	• **State:**
	  ◦ **Ack-rcvd**—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**—OSINCLP state completion time.
		extensive

Table 37: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

TAGCP	TAGCP information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  ◦ **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**—(M120 and M320 routers) One of the following values:
	  ◦ **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**—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 (LCP Connection Update Negotiation Successful)

```
user@host> show ppp interface extensive pp0.3221225489
```

```
  Session pp0.3221225489, Type: PPP, Phase: Network
   Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
               Magic-Number validation: enable
    LCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Last updated: 2020-02-11 15:06:10 PDT
      Negotiated options:
        Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492,
Peer MRU: 149
        Juniper Connection Status Message: 10m:xxxx
        Connection-Update-Requests: 1
     Authentication: Off
    IPCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local address: 198.51.100.30, Remote address: 203.0.113.9
      Negotiation mode: Passive
    IPV6CP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
       Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier:
```

```
 2001:db8::3a
      Negotiation mode: Passive
```

show ppp interface extensive (LCP Connection Update Negotiation Failed)

user@host> **show ppp interface extensive pp0.3221225489**

```
  Session pp0.3221225489, Type: PPP, Phase: Network
   Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
               Magic-Number validation: enable
    LCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492,
Peer MRU: 149
        Connection-Update-Requests: 1
     Authentication: Off
    IPCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local address: 198.51.100.30, Remote address: 203.0.113.9
      Negotiation mode: Passive
    IPV6CP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
       Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier:
 2001:db8::3a
      Negotiation mode: Passive
```

show ppp interface extensive (Inline Service Interface)

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 timout    :    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 917
show ppp statistics detail on page 918
show ppp statistics recovery (Safe to Restart) on page 919
show ppp statistics recovery (Unsafe to Restart) on page 919

Output Fields

Table 38 on page 913 lists the output fields for the **show ppp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 38: show ppp statistics Output Fields

Field Name	Field Description	Level of Output
Total sessions	Number of PPP sessions on an interface.	none **detail**
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 38: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type	Type of structure for which memory is allocated. • **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 38: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type (continued)	• **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 38: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type	If you specify the **memory** keyword, the following memory statistics are displayed for Ethernet interfaces on M120 and M320 routers. • **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**

Table 38: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
Failures	Number of failed allocations.	**detail**
Recovery state	State of PPP recovery after a GRES or restart: • 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
Interface address       64       0     0  65535           0         0
Destination profile     65       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
linkInterface           152       1    99  16384         152        0        0
pap                     256       1    99  16384         256        0        0
lcp                     272       1    99  16384         272        0        0
chap                    284       0     0  16384           0        0        0
eapBuffer              1464       0     0  16384           0        0        0
eap                     276       0     0  16384           0        0        0
authNone
networkInterface        220       1    99  16384         220        0        0
ipNcp                   256       1    99  16384         256        0        0
ipv6Ncp                 204       0     0  16384           0        0        0
osiNcp                  192       0     0  16384           0        0        0
mplsNcp                 188       0     0  16384           0        0        0
trace                  2052       0    16     16           0        0        0
Total                                                   1380        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 922

Output Fields

Table 39 on page 921 lists the output fields for the **show ppp summary** command. Output fields are listed in the approximate order in which they appear.

Table 39: 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
rlsq0.0        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.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the inline IP reassembly statistics for the Packet Forwarding Engines on one or more MPCs or Next Gen Services MX-SPC3 services card. 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 or MX-SPC3 services card.

fpc *fpc*—(Optional) Displays inline IP reassembly statistics for the specified MPC or MX-SPC3 services card.

> **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 | 616

List of Sample Output

show services inline ip-reassembly statistics fpc on page 927

Output Fields

Table 40 on page 924 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 40: show services inline ip-reassembly statistics Output Fields

Field Name	Field Description
FPC	MPC or MX-SPC3 services card slot number for which the statistics are displayed.
PFE	Packet Forwarding Engine on the MPC or MX-SPC3 services card 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	Total number of fragments received and the current rate of fragments received for inline IP reassembly. The following information is also displayed: • **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. **NOTE:** Current rate refers to the current number of fragments processed per second in the instant preceding the command's execution.
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.

Table 40: show services inline ip-reassembly statistics Output Fields *(continued)*

Field Name	Field Description
Approximate Packets Pending Reassembly	Approximate number of packets pending reassembly.
Fragments Dropped Reasons	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: • 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 **NOTE:** • These fields indicate *why* 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	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: • Fragment not found • Fragment not in sequence • ASIC errors **NOTE:** Current rate refers to the current number of reassembly errors processed per second in the instant preceding the command's execution.

Table 40: show services inline ip-reassembly statistics Output Fields *(continued)*

Field Name	Field Description
Aged out packets	Number of aged out packets and the current number of packets aged out per second in the instant preceding the command's execution. **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.
Total Fragments Successfully Reassembled	Number of fragments successfully reassembled and the current number of fragments reassembled per second in the instant preceding the command's execution.
Total Fragments Dropped	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: • 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 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

```
show services l2tp client
<client-name>
```

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

- show services l2tp session-limit-group | 949
- show services l2tp tunnel-group | 967
- L2TP Session Limits Overview | 207

List of Sample Output

show services l2tp client on page 930

show services l2tp client (Client Name) on page 930

Output Fields

Table 41 on page 929 lists the output fields for the **show services l2tp client** command. Output fields are listed in the approximate order in which they appear.

Table 41: show services l2tp client Output Fields

Field Name	Field Description
Client **Client Name**	Name of the client.
Tunnels	Number of tunnels in the tunnel group.

Table 41: show services l2tp client Output Fields *(continued)*

Field Name	Field Description
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

```
show services l2tp destination
<brief | detail | extensive>
<local-gateway gateway-address>
<peer-gateway gateway-address>
<statistics>
```

Release Information

Command introduced in Junos OS Release 10.4.

Description

Display information about L2TP tunnel destinations.

Options

brief | detail | extensive—(Optional) Display the specified level of information.

local-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified local gateway address.

peer-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified peer gateway address.

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**.

Required Privilege Level

view

RELATED DOCUMENTATION

- clear services l2tp destination | 844
- show services l2tp destination lockout | 936
- show services l2tp session | 938
- show services l2tp summary | 951
- show services l2tp tunnel | 959

List of Sample Output

show services l2tp destination on page 934

show services l2tp destination detail on page 934

show services l2tp destination extensive (LAC) on page 934
show services l2tp destination extensive (LNS) on page 935
show services l2tp destination statistics (LAC only on MX Series Routers) on page 935

Output Fields

Table 42 on page 932 lists the output fields for the **show services l2tp destination** command. Output fields are listed in the approximate order in which they appear.

Table 42: 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: • total • active • failed	All levels for total **extensive** for active and failed
Sessions	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed
State	Administrative state of the L2TP destination: • **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**

Table 42: show services l2tp destination Output Fields *(continued)*

Field Name	Field Description	Level of Output
Lockout State	Reachability state of the destination: • **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**
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

- clear services l2tp destination | 844
- request services l2tp destination unlock | 744
- show services l2tp destination | 931
- show services l2tp session | 938
- show services l2tp summary | 951
- show services l2tp tunnel | 959

List of Sample Output

show services l2tp destination lockout on page 937

Output Fields

Table 43 on page 936 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 43: show services l2tp destination lockout Output Fields

Field Name	Field Description
Destination	Name of the destination.

Table 43: show services l2tp destination lockout Output Fields *(continued)*

Field Name	Field Description
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 | **848**

List of Sample Output

show services l2tp session (LNS on M Series Routers) on page 944
show services l2tp session (LNS on MX Series Routers) on page 945
show services l2tp session (LAC) on page 945
show services l2tp session detail (LAC) on page 945
show services l2tp session extensive (LAC) on page 945
show services l2tp session extensive (LAC on MX Series Routers) on page 946
show services l2tp session extensive (LNS on M Series Routers) on page 946
show services l2tp session extensive (LNS on MX Series Routers) on page 947

show services l2tp session statistics (MX Series Routers) on page 948

Output Fields

Table 44 on page 940 lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

Table 44: show services l2tp session Output Fields

Field Name	Field Description	Level of Output
Interface	(LNS only) Name of an adaptive services interface.	All levels
Tunnel group	(LNS only) Name of a tunnel group.	All levels
Tunnel local ID	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	All levels
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
State	State of the L2TP session: • **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

Table 44: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Mode	(LNS) Mode of the interface representing the session: **shared** or **exclusive**. (LAC) Mode of the interface representing the session: **shared** or **dedicated**. Only **dedicated** is currently supported for the LAC.	**extensive**
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	**extensive**
Remote IP	IP address of remote endpoint of the PPP session.	**extensive**
Username	(LNS only) Name of the user logged in to the session.	All levels
Assigned IP address	(LNS only) IP address assigned to remote client.	**extensive**
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	**extensive**
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	**extensive**
Local MRU	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	**extensive**
Remote MRU	(LNS only) MRU setting of the remote device, in bytes.	**extensive**

Table 44: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tx speed	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. 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: • 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. For Junos OS Release 17.2 and Release 17.3, only the current (**update**) line speed can be displayed on MX Series routers. 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. Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.	**extensive**
Rx speed	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. 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: • 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. For Junos OS Release 17.2 and Release 17.3, only the current (**update**) line speed can be displayed on MX Series routers. 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. Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.	**extensive**

Table 44: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bearer type	Type of bearer enabled: • **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	Type of framing enabled: • **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**
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**

Table 44: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Up time	Length of time elapsed since the call became active, in hours, minutes, and seconds.	**extensive**
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	**extensive**
Statistics since	Date and time when collection of the following statistics began: • **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 timout**—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

- show services l2tp client | **929**
- show services l2tp tunnel-group | **967**
- L2TP Session Limits Overview | **207**

List of Sample Output

show services l2tp session-limit-group on page 950
show services l2tp session-limit-group (Limit Group Name) on page 950

Output Fields

Table 45 on page 949 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 45: 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.

Table 45: show services l2tp session-limit-group Output Fields *(continued)*

Field Name	Field Description
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       Maximum limit
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
```

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

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.

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.

statistics—(Optional) Display a summary of control packets and bytes transmitted and received.

Required Privilege Level

view

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

List of Sample Output

show services l2tp summary (LAC on M Series routers) on page 955
show services l2tp summary (LAC on MX Series routers) on page 956
show services l2tp summary (LNS on MX Series routers) on page 956
show services l2tp summary (LNS on M Series routers) on page 957
show services l2tp summary statistics (MX Series routers) on page 957

Output Fields

Table 46 on page 952 lists the output fields for the **show services l2tp summary** command. Output fields are listed in the approximate order in which they appear.

Table 46: 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.
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.

Table 46: show services l2tp summary Output Fields *(continued)*

Field Name	Field Description
Failover Protocol	When the state is enabled, the LAC operates in the default *failover-protocol-fall-back-to-silent-failover* 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	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: • **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**.
Tunnel assignment id	Format of the tunnel name. Format of the tunnel name, based on RADIUS attributes returned from the AAA server: • **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.

Table 46: show services l2tp summary Output Fields *(continued)*

Field Name	Field Description
Tunnel Tx Address Change	Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both: • **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 46: show services l2tp summary Output Fields *(continued)*

Field Name	Field Description
Access Line Information	State of LAC global configuration for forwarding subscriber line information to the LNS, **Enabled** or **Disabled**. Indicates active IP reassembly configured for the interface. Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for information it receives from the LAC.
IPv6 Services for LAC Sessions	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**.
Speed Updates	State of LAC global 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 can also be displayed on the LNS for updates it receives from the LAC.
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 963
show services l2tp tunnel detail (LAC) on page 963
show services l2tp tunnel detail (LAC on MX Series Routers) on page 963
show services l2tp tunnel detail (LNS on MX Series Routers) on page 964
show services l2tp tunnel extensive (LAC) on page 964
show services l2tp tunnel extensive (LNS on M Series Routers) on page 964
show services l2tp tunnel extensive (LNS on MX Series Routers) on page 965
show services l2tp tunnel statistics (MX Series Routers) on page 966

Output Fields

Table 47 on page 960 lists the output fields for the **show services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 47: 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.

Table 47: show services l2tp tunnel Output Fields *(continued)*

Field Name	Field Description
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.
State	State of the L2TP tunnel: • **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.

Table 47: show services l2tp tunnel Output Fields *(continued)*

Field Name	Field Description
Effective Peer Resync Mechanism	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.
Nas Port Method	NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS: • **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. The displayed limit for configured sessions is set to the lowest of the following configured session values for either LAC or LNS: • Global (chassis)—**set services l2tp tunnel** maximum-sessions ***number*** • Tunnel profile (individual tunnel)—**set access tunnel-profile** ***profile-name*** **tunnel** ***tunnel-id*** max-sessions ***number*****]** • RADIUS—Value of VSA 26–33, Tunnel-Max-Sessions For LNS only, the following configuration is also considered: • Host profile—**access profile l2tp-profile client default l2tp maximum-sessions-per-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.
Create time	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the when the call was created. 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.

Table 47: show services l2tp tunnel Output Fields *(continued)*

Field Name	Field Description
Statistics since	Date and time when collection of the following statistics began: • **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: 203.0.113.101:1701
     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: 192.0.2.3:1701
    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               0
```

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

```
show services l2tp tunnel-group
<group-name>
```

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

group-name—(Optional) Name of a tunnel group.

Required Privilege Level

view

RELATED DOCUMENTATION

- show services l2tp client | 929
- show services l2tp session-limit-group | 949
- L2TP Session Limits Overview | 207

List of Sample Output

show services l2tp tunnel-group on page 968

show services l2tp tunnel-group (Group Name) on page 968

Output Fields

Table 48 on page 967 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 48: 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

none—Display standard information for all L2TP switched tunnel destinations.

detail | extensive—(Optional) Display the specified level of information.

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**.

Required Privilege Level

view

RELATED DOCUMENTATION

- show services l2tp tunnel-switch session | 974
- show services l2tp tunnel-switch summary | 980
- show services l2tp tunnel-switch tunnel | 982

List of Sample Output

show services l2tp tunnel-switch destination on page 971
show services l2tp tunnel-switch destination detail on page 971
show services l2tp tunnel-switch destination extensive on page 972
show services l2tp tunnel-switch destination statistics on page 973

Output Fields

Table 49 on page 970 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 49: 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
Tunnels	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed
Sessions	Number of session connections for the destination in the following categories: • 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: • **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**

Table 49: show services l2tp tunnel-switch destination Output Fields *(continued)*

Field Name	Field Description	Level of Output
Lockout State	Reachability state of the destination: • **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**
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            192.0.2.3       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: 192.0.2.3
    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: 192.0.2.3
    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

none—Display standard information about all active L2TP switched tunnel sessions.

detail | extensive—(Optional) Display the specified level of output.

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**.

Additional Information

Required Privilege Level

view

RELATED DOCUMENTATION

- show services l2tp tunnel-switch destination | 969
- show services l2tp tunnel-switch summary | 980
- show services l2tp tunnel-switch tunnel | 982

List of Sample Output

show services l2tp tunnel-switch session on page 977
show services l2tp tunnel-switch session detail on page 977
show services l2tp tunnel-switch session extensive on page 978

Output Fields

Table 50 on page 975 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 50: 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
State	State of the L2TP session: • **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
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**

Table 50: show services l2tp tunnel-switch session Output Fields *(continued)*

Field Name	Field Description	Level of Output
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: • **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	Type of framing enabled: • **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**

Table 50: show services l2tp tunnel-switch session Output Fields *(continued)*

Field Name	Field Description	Level of Output
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: • **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: 192.0.2.3:1701
```

```
    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: 192.0.2.3:1701
    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 | 969
- show services l2tp tunnel-switch session | 974
- show services l2tp tunnel-switch tunnel | 982

List of Sample Output

show services l2tp tunnel-switch summary on page 981

Output Fields

Table 51 on page 980 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 51: show services l2tp tunnel-switch summary Output Fields

Field Name	Field Description
Tunnel switch profile name	Name of a tunnel switch profile.

Table 51: show services l2tp tunnel-switch summary Output Fields *(continued)*

Field Name	Field Description
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                                name
  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

none—Display standard information about all active L2TP tunnels.

detail | extensive—(Default) Display the specified level of output.

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**.

Additional Information

Required Privilege Level

view

RELATED DOCUMENTATION

- show services l2tp tunnel-switch destination | 969
- show services l2tp tunnel-switch session | 974
- show services l2tp tunnel-switch summary | 980

List of Sample Output

show services l2tp tunnel-switch tunnel on page 985
show services l2tp tunnel-switch tunnel detail on page 986
show services l2tp tunnel-switch tunnel extensive on page 986

Output Fields

Table 52 on page 983 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 52: show services l2tp tunnel-switch tunnel Output Fields

Field Name	Field Description	Level of Output
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.	none
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: • **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**

Table 52: show services l2tp tunnel-switch tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
Effective Peer Resync Mechanism	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • 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: • 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**

Table 52: show services l2tp tunnel-switch tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
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: • **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         192.0.2.3:1701               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: 192.0.2.3:1701
    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: 192.0.2.3:1701
    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

```
show services soft-gre tunnel
<brief | detail | extensive>
<interface interface-name>
<local-ip local-ip-address>
<remote-ip remote-ip-address>
<statistics>
<tunnel-group group-name>
```

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 (ps*n*) interface devices.

Options

none—Display standard information about all active dynamic GRE tunnels.

brief | detail | extensive—(Optional) Display the specified level of detail.

interface *interface-name*—(Optional) Display information for a specific pseudowire subscriber (ps*n*) interface.

local-ip *local-ip-address*—(Optional) Display information for a specific local or source IP address of the GRE tunnel endpoint (Wi-Fi access gateway side).

remote-ip *remote-ip-address*—(Optional) Display information for a specific remote IP address of the GRE tunnel endpoint.

statistics—(Optional) Display dynamic GRE tunnel statistics.

tunnel-group *group-name*—(Optional) Display information for a specific dynamic GRE tunnel group.

Required Privilege Level

view

RELATED DOCUMENTATION

Wi-Fi Access Gateway Deployment Model Overview | 360

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway | 361

Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation | 365
soft-gre | 777

List of Sample Output

show services soft-gre tunnel brief on page 990
show services soft-gre tunnel detail on page 990
show services soft-gre tunnel extensive on page 990

Output Fields

Table 53 on page 989 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 53: 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: *yyyy-mm-dd hh:mm:ss timezone*	**detail extensive**
Statistics since	Day of the week, date, time, and year when statistics for packets received and transmitted were recorded: *Dayofweek Month date hh:mm:ss yyyy*	**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**

Table 53: show services soft-gre tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
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     192.0.2.10               198.51.100.1                   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: 198.51.100.1
    Remote IP: 192.0.2.10
    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: 198.51.100.1
    Remote IP: 192.0.2.10
    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>
<agent-remote-identifier agent-remote-identifier>
<aggregation-interface-set-name interface-set-name>
<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 12.3R*x* 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 on MX Series routers.
accounting-statistics option added in Junos OS Release 15.1R3 and 17.4R1 on MX Series routers.
aggregation-interface-set-name option added in Junos OS Release 18.4R1 on MX Series routers.

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—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

Junos OS Release	Substring Support
Junos OS Release 13.3R1	You can specify a substring without a wildcard.
Starting in Junos OS Release 14.1R1	You must specify the complete ACI string; you cannot specify a wildcard.
Starting in Junos OS Release 15.1R7, 16.1R7, 16.2R3, 17.1R3, 17.2R3, 17.3R3, 17.4R2, 18.1R2, 18.2R1	You can specify a substring, but you must include the wildcard character at the end of the substring.

agent-remote-identifier—(Optional) Display all dynamic subscriber sessions whose ARI value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.

aggregation-interface-set-name *interface-set-name*—(Optional) Display summary information for the specified aggregation node interface set, including interface, VLAN ID, username and LS:RI.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dotlx**—Dotlx clients only.
- **essm**—ESSM clients only.

- **fixed-wireless-access**—Fixed wireless access 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. If the statement is not configured, a value of 0 is displayed for accounting statistics.

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 | **1040**
- *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 1007
show subscribers (IPv6) on page 1007
show subscribers (IPv4 and IPv6 Dual Stack) on page 1007
show subscribers (Single Session DHCP Dual Stack) on page 1008

show subscribers (Single Session DHCP Dual Stack detail) on page 1008
show subscribers (LNS on MX Series Routers) on page 1009
show subscribers (L2TP Switched Tunnels) on page 1009
show subscribers aggregation-interface-set-name on page 1009
show subscribers client-type dhcp detail on page 1009
show subscribers client-type dhcp detail (DHCPv6) on page 1010
show subscribers client-type dhcp extensive on page 1011
show subscribers client-type fixed-wireless-access on page 1012
show subscribers client-type fixed-wireless-access detail (Detail) on page 1012
show subscribers client-type vlan-oob detail on page 1013
show subscribers count on page 1013
show subscribers address detail (IPv6) on page 1013
show subscribers detail (IPv4) on page 1014
show subscribers detail (IPv6) on page 1015
show subscribers detail (pseudowire Interface for GRE Tunnel) on page 1015
show subscribers detail (IPv6 Static Demux Interface) on page 1016
show subscribers detail (L2TP LNS Subscribers on MX Series Routers) on page 1016
show subscribers detail (L2TP Switched Tunnels) on page 1016
show subscribers detail (Tunneled Subscriber) on page 1017
show subscribers detail (IPv4 and IPv6 Dual Stack) on page 1017
show subscribers detail (ACI Interface Set Session) on page 1019
show subscribers detail (PPPoE Subscriber Session with ACI Interface Set) on page 1019
show subscribers detail (Dynamic Profile Version Alias) on page 1020
show subscribers extensive on page 1020
show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes) on page 1021
show subscribers extensive (Passive Optical Network Circuit Interface Set) on page 1022
show subscribers extensive (DNS Addresses from Access Profile or Global Configuration) on page 1023
show subscribers extensive (DNS Addresses from RADIUS) on page 1024
show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration) on page 1024
show subscribers extensive (RPF Check Fail Filter) on page 1025
show subscribers extensive (L2TP LNS Subscribers on MX Series Routers) on page 1025
show subscribers extensive (IPv4 and IPv6 Dual Stack) on page 1026
show subscribers extensive (ADF Rules) on page 1027
show subscribers extensive (Effective Shaping-Rate) on page 1028
show subscribers extensive (PPPoE Subscriber Access Line Rates on page 1028
show subscribers extensive (Subscriber Session Using PCEF Profile) on page 1030
show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set) on page 1031
show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring) on page 1032
show subscribers id accounting-statistics on page 1033
show subscribers interface accounting-statistics on page 1033

show subscribers interface extensive on page 1034
show subscribers logical-system terse on page 1035
show subscribers physical-interface count on page 1035
show subscribers routing-instance inst1 count on page 1036
show subscribers stacked-vlan-id detail on page 1036
show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 1036
show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 1036
show subscribers user-name detail on page 1037
show subscribers vlan-id on page 1037
show subscribers vlan-id detail on page 1037
show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session) on page 1038
show subscribers address detail (Enhanced Subscriber Management) on page 1038
show subscribers extensive (Tenant Systems) on page 1039

Output Fields

Table 54 on page 997 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 54: 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 *tpid.vlan-id* 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, FWA, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
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.
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.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through NDRA.
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 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
Interface Set	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. The name of the interface set uses one of the following prefixes: • **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. **NOTE:** ACI interface sets are configured with the **agent-circuit-identifier** autoconfiguration stanza. ALI interface sets are configured with the **line-identity** autoconfiguration stanza. 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.
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**).

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
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).
PFE Flow ID	Forwarding flow identifier.
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	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. 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.
Agent Remote ID	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. 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.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
Aggregation Interface-set Name	Value of the $junos-aggregation-interface-set-name predefined variable; one of the following: • When the **hierarchical-access-network-detection** option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the ANCP Port Up message or PPPoE PADR IA tags begins with a **#** character, then the variable takes the value of the remainder of the string after the **#** character. • When the **hierarchical-access-network-detection** option is not configured, or if the sting does not begin with the **#** character, then the variable takes the value specified with the **predefined-variable-defaults** statement.
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.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
IPv4 Input Service Set	Input service set in access dynamic profile.
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.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
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.
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).

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
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	Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091): • **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. The value is stored in the subscriber session database.
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).
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.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
Local TEID-U	Tunnel endpoint identifier on the BNG for the GTP-U user plane tunnel to the eNodeB. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the **GTPU Tunnel Local IP address** value.
Local TEID-C	Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the MME. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the **GTPC Local IP address** value.
Remote TEID-U	Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB. A fully qualified remote TEID-U consists of this identifier and the **GTPU Tunnel Remote IP address** value.
Remote TEID-C	Tunnel endpoint identifier on the MME for the GTP-C control plane tunnel to the BNG. The identifier is allocated by the MME. A fully qualified remote TEID-C consists of this identifier and the **GTPC Remote IP address** value.
GTPU Tunnel Remote IP address	IP address of the S1-U interface on the eNodeB for the GTP-U tunnel endpoint. A fully qualified remote TEID-U consists of this address and the **Remote TEID-U** value.
GTPU Tunnel Local IP address	IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint. A fully qualified local TEID-U consists of this address and the **Local TEID-U** value
GTPC Remote IP address	IP address of the S11 interface on the MME for the GTP-C tunnel endpoint. A fully qualified remote TEID-C consists of this address and the **Remote TEID-C** value.
GTPC Local IP address	IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint. A fully qualified local TEID-C consists of this address and the **Local TEID-C** value.
Access Point Name	Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.

Table 54: show subscribers Output Fields *(continued)*

Field Name	Field Description
Tenant	Name of the tenant system. You can create multiple tenant system administrators for a tenant system with different permission levels based on your requirements.
Routing instance	Name of the routing instance. When a custom routing instance is created for a tenant system, all the interfaces defined in that tenant system are added to that routing instance.
Dynamic Profile Version Alias	Configured name for a specific variation of a base dynamic profile. IT's presence indicates that the profile configuration is different from that of the base profile. The value is conveyed to the RADIUS server during authentication in the Client-Profile-Name VSA (26–4874–174).

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
```

```
Interface               IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824     10                                    default:default
demux0.1073741824       203.0.113.10        WHOLESALER-CLIENT default:default
demux0.1073741825       203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826       203.0.113.3         RETAILER2-CLIENT  test1:retailer2
```

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    WHOLESALER-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
default:default
                    2001:db8:3ffe:0:4::/64
```

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 aggregation-interface-set-name

```
user@host> show subscribers aggregation-interface-set-name FRA*
```

```
Interface              IP Address/VLAN ID              User Name
LS:RI
ge-1/0/0.3221225472    50                                ancp
  default:isp1-subscriber
```

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 detail (DHCPv6)

`user@host>` **show subscribers client-type dhcp detail**

```
Type: DHCP
User Name: DEFAULTUSER
IPv6 Address: 2001:db8::2
IPv6 Prefix: 2001:db8:1::/64
Logical System: default
Routing Instance: default
Interface: demux0.3221225602
Interface type: Static
Underlying Interface: demux0.3221225602
Dynamic Profile Name: client-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 142
Session ID: 142
```

```
PFE Flow ID: 148
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2018-03-29 12:27:38 EDT
DHCP Options: len 56
00 08 00 02 00 00 00 01 00 0e 00 01 00 01 22 4f d0 33 00 11
01 00 00 01 00 03 00 0c 00 00 00 0a 00 04 9d 40 00 07 62 00
00 19 00 0c 00 00 00 0b 00 04 9d 40 00 07 62 00
Server DHCPV6 Options: len 94
00 0a 00 06 11 22 33 44 55 66 00 11 00 09 00 00 0c 4c 00 02
00 01 aa 00 11 00 20 00 00 0a 4c 00 02 00 02 32 33 00 03 00
03 34 35 36 00 05 00 06 31 32 33 34 35 36 00 06 00 01 31 00
11 00 09 00 00 0b 4c 00 02 00 01 bb 00 11 00 12 00 00 0d e9
00 01 00 03 aa bb cc 00 02 00 03 dd ee cc
DHCPV6 Header: len 4
01 fc e4 96
```

show subscribers client-type dhcp extensive

user@host> **show subscribers client-type dhcp extensive**

```
Type: DHCP
User Name: user
IP Address: 192.0.2.4
IP Netmask: 255.0.0.0
IPv6 Address: 2001:db8:3::103
IPv6 Prefix: 2001:db8::/68
Domain name server inet6: 2001:db8:1 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: 00:00:5e:00:53:01
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
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 fixed-wireless-access

`user@host>` **show subscribers client-type fixed-wireless-access**

```
Interface             IP Address/VLAN ID            User Name
LS:RI
ps1.3221225472        192.0.2.10                    505024101215074
default:default
ps1.3221225473        192.0.2.11                    505024101215075
default:default
```

show subscribers client-type fixed-wireless-access detail (Detail)

`user@host>` **show subscribers client-type fixed-wireless-access detail**

```
Type: FWA
User Name: 505024101215074
IP Address: 192.0.2.10
IP Netmask: 255.255.0.0
Interface: ps1.3221225472
Interface type: Dynamic
Dynamic Profile Name: fwa-profile
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 11
Login Time: 2019-04-10 14:10:12 PDT
Local TEID-U: 1
Local TEID-C: 1
```

```
Remote TEID-U: 2000000
Remote TEID-C: 1000000
GTPU Tunnel Remote IP Address: 203.0.113.1.3
GTPU Tunnel Local IP Address: 203.0.113.2.5
GTPC Remote IP Address: 203.0.113.1.2
GTPC Local IP Address: 203.0.113.1.1
Access Point Name: user21
```

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: ifl3720
Agent Remote ID: ifl3720
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 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      192.0.2.2
ps0.3221225485      192.0.2.3
demux0.3221225486    1                                         default:default

demux0.3221225487    1                                         default:default

demux0.3221225488   198.51.0.1                                 default:default

demux0.3221225489   198.51.0.2                                 default:default
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
```

```
Type: STATIC-INTERFACE
User Name: user@example.com
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.com
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 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 detail (Dynamic Profile Version Alias)

user@host> **show subscribers detail**

```
Type: PPPoE
User Name: DEFAULTUSER
IP Address: 192.0.2.21
IP Netmask: 255.255.255.255
IPv6 Address: 2001:db8::17
Logical System: default
Routing Instance: default
Interface: pp0.3221225720
Interface type: Dynamic
Underlying Interface: demux0.3221225719
Dynamic Profile Name: pppoe-client-profile
Dynamic Profile Version Alias: profile-version1a
MAC Address: 00:00:5E:00:53:38
State: Active
Radius Accounting ID: 288
Session ID: 288
PFE Flow ID: 344
VLAN Id: 1
Login Time: 2019-09-23 10:40:56 IST
```

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 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes)

`user@host>` **show subscribers extensive**

```
Type: VLAN-OOB
User Name: ancp
Logical System: default
Routing Instance: isp1-subscriber
Interface: ge-1/0/0.3221225472
Interface type: Dynamic
Interface Set: FRA-DPU-C-100
Underlying Interface: ge-1/0/0
Core IFL Name: ge-1/0/4.0
Dynamic Profile Name: Prof_L2BSA
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 13
VLAN Id: 50
VLAN Map Id: 20
Inner VLAN Map Id: 1
Inner VLAN Tag Protocol Id: 0x88a8
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:42 EDT
Accounting interval: 72000
Dynamic configuration:
  junos-cos-scheduler-map: 100m
  junos-inner-vlan-tag-protocol-id: 0x88a8
  junos-vlan-map-id: 20

Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
```

```
Routing Instance: default
Interface: pp0.3221225474
Interface type: Dynamic
Interface Set: ge-1/0/0
Underlying Interface: demux0.3221225473
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:10:94:00:00:03
State: Active
Radius Accounting ID: 3
Session ID: 3
PFE Flow ID: 16
Stacked VLAN Id: 50
VLAN Id: 7
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:45 EDT
IP Address Pool: pool-1
Accounting interval: 72000
DSL type: G.fast
Frame/cell mode: Frame
Overhead accounting bytes: 10
Actual upstream data rate: 100000 kbps
Actual downstream data rate: 200000 kbps
Calculated downstream data rate: 180000 kbps
Calculated upstream data rate: 90000 kbps
Adjusted upstream data rate: 80000 kbps
Adjusted downstream data rate: 160000 kbps
DSL Line Attributes
  Agent Circuit ID: circuit 201
  Agent Remote ID: remote-id
  Actual upstream data rate: 100000
  Actual downstream data rate: 200000
  DSL type: G.fast
  Access Aggregation Circuit ID: #FRA-DPU-C-100
  Attribute type: 0xAA, Attribute length: 4
    198 51 100 78
```

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: otl01.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: user :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|otl01.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.com
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: 198.51.100.1
 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: 00:00:5e:00:53:02
 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-88a8-400LL1300VO
  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-400LL1300VO-ge-1/2/4.3221225472-in
  IFL Output Filter Name: L2BSA-CP-400LL1300VO-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: 192.0.2.26
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: 00:00:5e:00:53:01
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 extensive (Tenant Systems)

`user@host:TSYS1>` **show subscribers extensive**

```
Type: XAUTH
User Name: userX
+ Tenant: TSYS1
 Routing Instance: TSYS1-ri
IP Address: 192.0.2.0
IP Netmask: 203.0.113.0
Primary DNS Address: 198.51.100.0
Secondary DNS Address: 198.51.100.1
Dynamic Profile Name: radius
State: Active
Session ID: 1
Login Time: 2018-09-18 13:49:00 PDT
```

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 | 992

List of Sample Output

show subscribers summary on page 1043
show subscribers summary all on page 1044
show subscribers summary physical-interface on page 1044
show subscribers summary physical-interface pic on page 1045
show subscribers summary physical-interface port on page 1045
show subscribers summary physical-interface slot on page 1046
show subscribers summary pic on page 1046
show subscribers summary pic (Aggregated Ethernet Interfaces) on page 1046
show subscribers summary port on page 1046
show subscribers summary port (Pseudowire Interfaces) on page 1047
show subscribers summary port extensive on page 1047
show subscribers summary slot on page 1047
show subscribers summary terse on page 1048

Output Fields

Table 55 on page 1042 lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 55: 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: • 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**

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.
	For aggregated Ethernet interfaces, the output of the **summary (pic | port | slot)** options prefixes the interface name with ae0:.
	For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and redundant logical tunnel (RLT) anchor interface. For example:
	`ps0: lt-2/1/0`
	`ps1: rlt0: lt-4/0/0`
		All levels

Table 55: show subscribers summary Output Fields *(continued)*

Field Name	Field Description	Level of Output
Count	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. 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. 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.	**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 *tpid.vlan-id*	**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   188
 Terminating    2
 Terminated    1

 TOTAL        191
```

```
Subscribers by Client Type
 DHCP     107
 PPP       76
 VLAN         8
 VLAN-OOB       2
 TOTAL        196
```

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                                   default:default
demux0.1073741824     203.0.113.10         WHOLESALER-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>
<dvlan>
<l2tp>
<ppp>
<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 protocols.

dhcp—(Optional) Display DHCP packet statistics.

dvlan—(Optional) Display DVLAN packet statistics.

l2tp—(Optional) Display L2TP packet statistics.

ppp—(Optional) Display PPP packet statistics.

pppoe—(Optional) Display PPPoE packet statistics.

Required Privilege Level

view

RELATED DOCUMENTATION

Understanding Dropped Packets and Untransmitted Traffic Using show Commands

List of Sample Output

show system subscriber-management statistics all on page 1051
show system subscriber-management statistics dhcp on page 1052
show system subscriber-management statistics dhcp extensive on page 1052
show system subscriber-management statistics pppoe on page 1055
show system subscriber-management statistics extensive on page 1055
show system subscriber-management statistics ppp (LCP Vendor-Specific Counters) on page 1056

Output Fields

Table 56 on page 1050 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 56: 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.
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 .
Layer 3 Statistics	Statistics for Layer 3 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, hybrid access, 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 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 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 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 statistics ppp (LCP Vendor-Specific Counters)

`user@host>` **show system subscriber-management statistics ppp**

```
Session Manager started @ Thu Feb 11 00:37:43 2020
Session Manager cleared @ Thu Feb 11 00:37:43 2020
-------------------------------------------------------------
                    Packet Statistics
-------------------------------------------------------------
```

```
I/O Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                         : 486783
    Tx Statistics
        packets                         : 144
  Layer 3 Statistics
    Rx Statistics
        packets                         : 8
    Tx Statistics
        packets                         : 0
PPP Statistics:
---------------------------------------------------------------
    Rx Statistics
        network packets                 : 123
        plugin packets                  : 123
        lcp config requests             : 18
        lcp config acks                 : 18
        lcp conf nacks                  : 8
        lcp conf rejects                : 6
        lcp termination requests        : 4
        lcp termination acks            : 13
        lcp code rejects                   : 2
        lcp vendor-specific acks       : 10
        pap requests                    : 8
        ipcp requests                   : 27
        ipcp acks                       : 9
        ipv6cp requests                 : 11
        ipv6cp acks                     : 1
    Tx Statistics
        packets                         : 101
        lcp config requests             : 32
        lcp config acks                 : 18
        lcp termination requests        : 13
        lcp termination acks            : 4
        lcp vendor-specific requests   : 10
        pap acks                        : 8
        ipcp requests                   : 9
        ipcp acks                       : 5
        ipcp nacks                      : 9
        ipv6cp requests                 : 1
        ipv6cp acks                     : 1
        ipv6cp nacks                    : 1
NET Statistics:
```

```
----------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                        : 8
        router solicitations            : 8
    Tx Statistics
        packets:                        : 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

show database-replication statistics

show database-replication summary

List of Sample Output

show system subscriber-management summary on page 1061
show system subscriber-management summary (Enhanced Subscriber Management) on page 1062
show system subscriber-management summary (DRAM Size Mismatch Error) on page 1062

Output Fields

Table 57 on page 1060 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 57: show system subscriber-management summary Output Fields

Field Name	Field Description
Graceful Restart	State of graceful Routing Engine switchover (GRES): • **Enabled** • **Disabled** (Enhanced subscriber management for MX Series routers) The name of this field is **Graceful Switchover**.
Mastership	State of the Routing Engine: • **Master** • **Standby**
Database	State of the subscriber management database: • **Available** • **Init** • **Not-available**
Standby	(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine: • **Connected**—Connected but not synchronized • **Disconnected**—Not connected • **Resync (*nn*%)**—Connected and *nn* percent synchronized with the master Routing Engine • **Synchronized**—Synchronized with the master Routing Engine
Disconnection Reason	Reason why both Routing Engines are disconnected when there is a DRAM mismatch. • Master/Standy RE DRAM Size Mismatch—Displayed when the amount of memory is different on the master and standby Routing Engines.

Table 57: show system subscriber-management summary Output Fields *(continued)*

Field Name	Field Description
Chassisd ISSU State	State of unified ISSU chassis daemon: • **ABORT** • **DAEMON_ISSU_PREPARE** • **DAEMON_ISSU_PREPARE_DONE** • **DAEMON_SWITCHOVER_PREPARE** • **DAEMON_SWITCHOVER_PREPARE_DONE** • **FRU_ISSU** • **FRU_ISSU_DONE** • **IDLE** • **UNKNOWN**
ISSU State	State of unified ISSU aggregate daemon: • **ABORT** • **IDLE** • **PREPARE** • **READY** • **SWITCHOVER_PREPARE** • **SWITCHOVER_READY** • **UNKNOWN**
ISSU Wait	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.

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
```

show system subscriber-management summary (DRAM Size Mismatch Error)

```
user@host> show system subscriber-management summary
```

```
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Standby              Disconnected
<emphasis>    Disconnection Reason    Master/Standy RE DRAM Size Mismatch</emphasis
>
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

test services l2tp tunnel

Syntax

```
test services l2tp tunnel user user-name
<password user-password>
<tunnel-name name>
```

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

user *user-name*—Name of the user under test. You must use an existing configured username, although it can be created solely for testing a tunnel configuration.

password *user-password*—(Optional) Authentication password for the specified user. If you omit this option, the test generates a dummy password—*testpass*—for the user.

tunnel-name *name*—(Optional) Name of a tunnel to test.

Required Privilege Level

view

RELATED DOCUMENTATION

Testing L2TP Tunnel Configurations from the LAC | 222

List of Sample Output

test services l2tp tunnel user (User authentication fails) on page 1064
test services l2tp tunnel user (Multiple tunnels tested) on page 1064
test services l2tp tunnel user tunnel-name (Specific tunnel tested) on page 1064

Output Fields

Table 58 on page 1064 lists the output fields for the **test services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 58: 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.
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    192.0.2.3        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    192.0.2.3        default         default         Up
```