

Junos® OS

Broadband Subscriber Sessions User Guide

Published
2026-01-09

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 Sessions User Guide
Copyright © 2026 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 This Guide | xxiii

1

AAA for Subscriber Management

AAA for Subscriber Management | 2

AAA Service Framework Overview | 2

Standard and Vendor-Specific RADIUS Attributes | 3

RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework | 4

RADIUS IETF Attributes Supported by the AAA Service Framework | 4

Juniper Networks VSAs Supported by the AAA Service Framework | 19

AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs | 599

AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs | 614

DSL Forum Vendor-Specific Attributes | 626

DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS | 637

RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses | 642

Support for Cisco Systems VSAs | 643

Subscriber Management RADIUS Dictionary Files | 643

Interface Text Descriptions for Inclusion in RADIUS Attributes | 643

Subscriber Management RADIUS Dictionary Files | 645

RADIUS for Subscriber Management | 646

RADIUS Servers and Parameters for Subscriber Access | 646

RADIUS Authentication and Accounting Server Definition | 647

Configuring Options that Apply to All RADIUS Servers | 650

Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | 652

Configuring Access Profile Options for Interactions with RADIUS Servers | 653

Configuring a Calling-Station-ID with Additional Options | 660

Filtering RADIUS Attributes and VSAs from RADIUS Messages | 665

Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers | 668

Interface Description Storage and Reporting Overview | 668

Interface Description Storage and Reporting Configuration | 673

Session Options for Subscriber Access | 674

Understanding Session Options for Subscriber Access | 674

Subscriber Session Timeout Options | 681

Limiting the Number of Active Sessions per Username and Access Profile | 683

Configuring Username Modification for Subscriber Sessions | 683

Removing Inactive Dynamic Subscriber VLANs | 687

RADIUS NAS Port Attributes and Options | 689

Manual Configuration of the NAS-Port-ID RADIUS Attribute | 689

Configuring a NAS-Port-ID with Additional Options | 691

Configuring the Order in Which Optional Values Appear in the NAS-Port-ID | 692

Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers | 694

RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview | 695

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 697

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 698

Manual Configuration of the NAS-Port-Type RADIUS Attribute | 699

Configuring the RADIUS NAS-Port-Type per Physical Interface | 702

Configuring the RADIUS NAS-Port-Type per VLAN | 703

Configuring the RADIUS NAS-Port-Type per Stacked VLAN | 705

Configuring the RADIUS NAS-Port Extended Format per Physical Interface | 707

Configuring the RADIUS NAS-Port Extended Format per VLAN | 708

Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN | 710

Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces | 712

RADIUS Logical Line Identification | 714

RADIUS Logical Line Identifier (LLID) Overview | 715

RADIUS Attributes for LLID Preauthentication Requests | 716

Configuring Logical Line Identification (LLID) Preauthentication | 718

Configuring a Port and Password for LLID Preauthentication Requests | 719

Verifying and Managing LLID Preauthentication Configuration | 721

RADIUS Authentication and Accounting Basic Configuration | 721

Configuring Authentication and Accounting Parameters for Subscriber Access | 722

Specifying the Authentication and Accounting Methods for Subscriber Access | 722

Specifying RADIUS Authentication and Accounting Servers for Subscriber Access | 723

| [Configuring Local Authentication and Authorization for Subscribers](#) | **724**

[RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers](#) | **727**

[Configuring RADIUS Reauthentication for DHCP Subscribers](#) | **738**

[RADIUS Accounting for Subscriber Access](#) | **741**

| [RADIUS Accounting Statistics for Subscriber Access Overview](#) | **742**

| [RADIUS Acct-On and Acct-Off Messages](#) | **744**

| [Configuring Per-Subscriber Session Accounting](#) | **744**

| [Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI](#) | **748**

| [Understanding RADIUS Accounting Duplicate Reporting](#) | **749**

| [Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting](#) | **751**

| [Configuring Per-Service Session Accounting](#) | **752**

| [Processing Cisco VSAs in RADIUS Messages for Service Provisioning](#) | **754**

| [Configuring Service Packet Counting for Volume Statistics](#) | **756**

| [Configuring Service Accounting](#) | **757**

| [Preservation of RADIUS Accounting Information During an Accounting Server Outage](#) | **759**

| [Configuring Back-up Options for RADIUS Accounting](#) | **762**

| [Forcing the Router to Contact the Accounting Server Immediately](#) | **763**

| [Monitoring Pending RADIUS Accounting Stop Messages](#) | **764**

| [Suspending RADIUS Accounting and Baseline Accounting Statistics Overview](#) | **766**

| [Configuring RADIUS Accounting Suspension and Baseline Accounting Statistics](#) | **770**

[Verifying and Managing Subscriber AAA Information](#) | **772**

[Session Termination Causes and RADIUS Termination Cause Codes](#) | **774**

| [Understanding Session Termination Causes and RADIUS Termination Cause Codes](#) | **775**

| [Mapping Session Termination Causes to Custom Termination Cause Codes](#) | **777**

[AAA Termination Causes and Code Values](#) | **779**

[DHCP Termination Causes and Code Values](#) | **781**

[L2TP Termination Causes and Code Values](#) | **783**

[PPP Termination Causes and Code Values](#) | **809**

[VLAN Termination Causes and Code Values](#) | **822**

[Domain Maps for Subscriber Management](#) | **825**

[Mapping Subscriber Domains to Access and Session Options](#) | **825**

Domain Mapping Overview	826
Configuring a Domain Map	830
Configuring a Wildcard Domain Map	832
Specifying an Access Profile in a Domain Map	833
Specifying an Address Pool in a Domain Map	834
Specifying a Dynamic Profile in a Domain Map	834
Specifying an AAA Logical System/Routing Instance in a Domain Map	835
Specifying a Target Logical System/Routing Instance in a Domain Map	836
Specifying a Tunnel Profile in a Domain Map	837
Specifying a Tunnel Switch Profile in a Domain Map	837
Configuring Domain and Realm Name Usage for Domain Maps	838
Specifying Domain and Realm Name Delimiters	839
Specifying the Parsing Order for Domain and Realm Names	840
Specifying the Parsing Direction for Domain and Realm Names	840
Enabling Domain Name Stripping	842
Changing the Username and Password to Simplify Off-Chassis Provisioning	842

Verifying Domain Maps | 844

Testing and Troubleshooting AAA | 846

AAA Testing and Troubleshooting | 846

AAA Configuration Testing and Troubleshooting	846
Testing a Subscriber AAA Configuration	847

Tracing General Authentication Service (authd) Events for Troubleshooting | 854

Configuring the General Authentication Service Trace Log Filename	855
Configuring the Number and Size of General Authentication Service Log Files	855
Configuring Access to the General Authentication Service Log File	856
Configuring a Regular Expression for General Authentication Service Messages to Be Logged	856
Configuring Subscriber Filtering for General Authentication Service Tracing	857
Configuring the General Authentication Service Tracing Flags	858

RADIUS Dictionary Files | 859

Junos OS Release 15.1 Subscriber Management RADIUS Dictionary [DCT] | 859

Junos OS Release 16.1 Subscriber Management RADIUS Dictionary [DCT] | 881

Junos OS Release 16.2 Subscriber Management RADIUS Dictionary [DCT] | 902

Junos OS Release 17.1 Subscriber Management RADIUS Dictionary [DCT] | 924

Junos OS Release 17.4 Subscriber Management RADIUS Dictionary [DCT] | 946

Junos OS Release 18.2 Subscriber Management RADIUS Dictionary [DCT] | 968

Junos OS Release 18.4 Subscriber Management RADIUS Dictionary [DCT] | 990

DHCP and DHCPv6 for Subscriber Management

DHCP for Subscriber Management | 1014

DHCP Overview | 1015

Understanding Differences Between Legacy DHCP and Extended DHCP | 1015

Extended DHCP Relay Agent Overview | 1019

DHCP Relay Proxy Overview | 1021

Minimum DHCP Relay Agent Configuration | 1023

Example: DHCP Relay Agent Configuration with Multiple Clients and Servers | 1027

DHCP Access Profiles for Subscriber Authentication and Accounting Parameters | 1029

Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview | 1029

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | 1030

Attaching an Access Profile to All DHCP Subscriber or All DHCP Client Interfaces | 1030

Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients | 1031

Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces | 1032

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033

Overriding the Default DHCP Local Server Configuration Settings | 1034

Overriding the Default DHCP Relay Configuration Settings | 1036

DHCP Behavior When Renegotiating While in Bound State | 1039

Sending Release Messages When Clients Are Deleted | 1040

Disabling Automatic Binding of Stray DHCP Requests | 1041

Enabling DHCP Relay Proxy Mode | 1043

Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent | 1044

Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address | 1044

Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall | 1045

Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets | 1045

Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally | 1046

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | 1047

- Load Balancing DHCP Local Servers by Delaying Responses to Clients | 1047

- Configuring a Delay in Local Server Response to DHCP Discover and Solicit Messages | 1048

DHCP Options and Selective Traffic Processing | 1051

- DHCP Options and Selective Traffic Processing Overview | 1052

- Using DHCP Option Information to Selectively Process DHCP Client Traffic | 1054

- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings | 1055

- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings | 1055

 - Requirements | 1056

 - Overview | 1056

 - Configuration | 1057

 - Verification | 1059

- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing | 1061

 - Requirements | 1061

 - Overview | 1062

 - Configuration | 1062

 - Verification | 1065

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | 1066

- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | 1067

 - Client-Side Support | 1069

 - Server-Side Support | 1069

 - DHCP Local Server Support | 1070

DHCP-Initiated Service Change Based on Remote ID | 1071

- Configuring DHCP-Initiated Service Change Based on Remote ID | 1072

- DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 1074

Using DHCP Option 82 Information | 1078

- Using DHCP Relay Agent Option 82 Information | 1078

 - Configuring Option 82 Information | 1079

 - Overriding Option 82 Information | 1082

 - Including a Prefix in DHCP Options | 1083

 - Including a Textual Description in DHCP Options | 1086

Enable Processing of Untrusted Packets So Option 82 Information Can Be Used | **1088**

Extracting an Option 82 or Option 37 Substring to Create an Interface Set | **1089**

Default Services for DHCP Subscribers | **1091**

Default Subscriber Service Overview | **1091**

Configuring a Default Subscriber Service | **1092**

DHCP Client Attribute and Address Assignment | **1093**

DHCP Attributes Overview | **1094**

Attributes That Can Be Applied to DHCP Clients | **1095**

Configuring DHCP Attributes for All Clients or a Group of Clients | **1098**

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | **1100**

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **1101**

Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option | **1103**

Specifying the Subnet for DHCP Client Address Assignment | **1104**

DHCP Local Server Handling of Client Information Request Messages | **1104**

Enabling Processing of Client Information Requests | **1105**

DNS Address Assignment Precedence | **1106**

Example: Extended DHCP Local Server Configuration with Optional Pool Matching | **1106**

DHCP Lease Times for IP Addresses | **1108**

DHCP Lease Timers | **1108**

DHCP Lease-Time Validation Overview | **1109**

Configuring a DHCP Lease-Time Threshold | **1111**

DHCP Asymmetric Leasing Overview | **1112**

Configuring DHCP Asymmetric Leasing | **1113**

DHCP Leasequery Methods | **1116**

Benefits of DHCP Leasequery | **1117**

DHCP Individual Leasequery | **1117**

DHCP Bulk Leasequery | **1121**

DHCP Active Leasequery | **1127**

Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations | **1138**

Configuring and Using DHCP Individual Leasequery | **1139**

Configuring and Using DHCP Bulk Leasequery | **1141**

Configuring and Using DHCP Active Leasequery | **1145**

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | **1152**

Verifying and Managing DHCP Individual and Bulk Leasequery Configurations | **1157**

Verifying and Managing DHCP Active Leasequery Operations | **1158**

DHCP Client Authentication With An External AAA Authentication Service | **1161**

Specifying Authentication Support | **1161**

Creating Unique Usernames for DHCP Clients | **1162**

Example-Configuring DHCP with External Authentication Server | **1164**

Receiving DHCP Options From a RADIUS Server | **1166**

Centrally Configure DHCP Options on a RADIUS Server | **1166**

Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview | **1170**

Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6
Subscribers | **1174**

Monitoring DHCP Options Configured on RADIUS Servers | **1176**

Common DHCP Configuration for Interface Groups and Server Groups | **1179**

Grouping Interfaces with Common DHCP Configurations | **1179**

Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces | **1182**

Configuring Group-Specific DHCP Local Server Options | **1183**

Configuring Group-Specific DHCP Relay Options | **1184**

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to
Named Server Groups | **1185**

Number of DHCP Clients Per Interface | **1189**

Specifying the Maximum Number of DHCP Clients Per Interface | **1189**

Allowing Only One DHCP Client Per Interface | **1191**

Maintaining DHCP Subscribers During Interface Delete Events | **1192**

Maintaining Subscribers During Interface Delete Events | **1193**

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | **1194**

Configuring an ACX Series DHCP Local Server to Preserve Subscriber Binding Information | **1194**

Verifying and Managing DHCP Subscriber Binding During Interface Delete Events | **1196**

Dynamic Reconfiguration of Clients From a DHCP Local Server | **1197**

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | **1198**

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | **1201**

Configuring Dynamic Reconfiguration Attempts for DHCP Clients | **1202**

Configuring Deletion of the Client When Dynamic Reconfiguration Fails | **1203**

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | **1204**

Configuring a Token for DHCP Local Server Authentication | **1204**

Understanding Deferred NACK on DHCP Reconfigure Abort | **1206**

Conserving IP Addresses Using DHCP Auto Logout | **1208**

DHCP Auto Logout Overview | **1208**

Automatically Logging Out DHCP Clients | **1210**

How DHCP Relay Agent Uses Option 82 for Auto Logout | **1211**

DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers | **1212**

Automatically Logging Out DHCPv6 Clients | **1213**

DHCP Short Cycle Protection | **1215**

DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions | **1215**

Configuring DHCP Short-Cycle Protection | **1218**

Verifying and Managing DHCP Short-Cycle Protection | **1222**

DHCP Monitoring and Management | **1224**

Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings | **1225**

Viewing and Clearing DHCP Bindings | **1226**

Monitoring DHCP Relay Server Responsiveness | **1228**

Verifying DHCP Server Binding and Server Statistics | **1229**

Verifying and Managing DHCP Relay Configuration | **1230**

Tracing Extended DHCP Operations | **1231**

Configuring the Extended DHCP Log Filename | **1233**

Configuring the Number and Size of Extended DHCP Log Files | **1233**

Configuring Access to the Extended DHCP Log File | **1234**

Configuring a Regular Expression for Extended DHCP Messages to Be Logged | **1235**

Configuring the Extended DHCP Tracing Flags | **1236**

Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged | **1236**

Tracing Extended DHCP Operations for Specific Interfaces | **1237**

3

IPv6 for Subscriber Management

IPv6 for Subscriber Management | **1241**

Introduction to IPv6 Addresses | **1241**

IPv6 Notation | **1242**

IPv6 Prefixes | **1242**

IPv6 Address Types | **1243**

Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245

Basic Architecture of a Subscriber Access Dual-Stack Network | 1245

Terms Used in IPv6 Subscriber Management Documentation | 1246

IPv6 Addressing Requirements for a Subscriber Access Network | 1248

IPv6 WAN Link Addressing with NDRA | 1250

Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250

IPv6 Neighbor Discovery Protocol Overview | 1252

Dynamic Router Advertisement Configuration Overview | 1253

Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | 1253

Methods for Obtaining IPv6 Prefixes for NDRA | 1255

Duplicate Prefix Protection for NDRA | 1256

IPv6 WAN Link Addressing with DHCPv6 IA_NA | 1257

Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA | 1258

Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | 1258

Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 1259

Using DHCPv6 Prefix Delegation Overview | 1260

Using a Delegated Prefix on the CPE Loopback Interface | 1261

DHCPv6 Prefix Delegation over PPPoE | 1261

Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation | 1262

Subscriber Delegated Prefix Preservation | 1263

DHCPv6 Prefix Exclusion | 1264

Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation | 1267

Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | 1267

WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1269

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269

DHCPv6 Options in a DHCPv6 Multiple Address Environment | 1271

Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA | 1272

Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface | 1273

Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | 1274

Requirements | 1274

Overview | 1274

Configuration | 1276

Verification | 1298

Designs for IPv6 Addressing in a Subscriber Access Network | 1305

- Selecting the Type of Addressing Used on the CPE | 1305
- Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link | 1306
- Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | 1306
- Selecting the Method of Obtaining IPv6 Prefixes | 1307
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1308
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | 1309
- Design 3: IPv6 Addressing with NDRA | 1310
- Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | 1311

Dual-Stack Access Models in a DHCP Network | 1312

- IPv4 and IPv6 Dual Stack in a DHCP Access Network | 1313
- AAA Service Framework in a Dual Stack over a DHCP Access Network | 1314
- Dual-Stack Interface Stack in a DHCP Wholesale Network | 1315
- Single-Session DHCP Dual-Stack Overview | 1316
- Configuring Single-Session DHCP Dual-Stack Support | 1320
- Verifying and Managing DHCP Dual-Stack Configuration | 1323

Dual-Stack Access Models in a PPPoE Network | 1324

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1325
- Shared IPv4 and IPv6 Service Sessions on PPP Access Networks | 1328
- AAA Service Framework in a Dual Stack over a PPPoE Access Network | 1328
- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers | 1331
- Accounting Messages for PPPoE Using NDRA Prefixes | 1332
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes | 1339
- Suppressing Accounting Information That Comes from AAA | 1349
- Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address | 1350

Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1351

- Best Practice: Static PPPoE Interfaces with NDRA | 1351
- Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network | 1352
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | 1353
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 | 1353
- Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles | 1354
- Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | 1355

Dual Stack for PPPoE Access Networks Using DHCP | 1356

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | **1356**

Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network | **1357**

Dual Stack for PPPoE Access Networks Using NDRA | **1360**

Configuring a Static PPPoE Logical Interface for NDRA | **1361**

Configuring an Address-Assignment Pool Used for Router Advertisements | **1362**

Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | **1363**

Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | **1364**

Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | **1364**

Requirements | **1365**

Overview | **1365**

Configuration | **1366**

Verification | **1384**

Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | **1390**

Requirements | **1390**

Overview | **1390**

Configuration | **1392**

Verification | **1414**

IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services | **1422**

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | **1422**

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | **1424**

Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1425**

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1426**

On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview | **1426**

On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview | **1429**

IPCP Negotiation with Optional Peer IP Address | **1431**

How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled | **1432**

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1433**

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1433**

Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1434**

Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes | **1434**

Enabling IPv4 Release Control VSA (26–164) in RADIUS Messages | 1435

Dual Stack Subscribers Monitoring and Management | 1436

Monitoring Active Subscriber Sessions | 1436

Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance | 1437

Monitoring Dynamic Subscriber Sessions | 1438

Monitoring Address Pools Used for Subscribers | 1440

Monitoring Specific Subscriber Sessions | 1442

Monitoring the Status of the PPPoE Logical Interface | 1443

Monitoring Service Sessions for Subscribers | 1444

Monitoring PPP Options Negotiated with the Remote Peer | 1445

Monitoring the RADIUS Attribute Used for NDRA | 1447

4

DHCPv6 for Subscriber Management

DHCPv6 Local Server | 1449

DHCPv6 Local Server Overview | 1449

Configuring DHCPv6 Rapid Commit (MX Series, EX Series) | 1450

Preventing Binding of Clients That Do Not Support Reconfigure Messages | 1451

Configuring the DUID Type Supported by DHCPv6 Servers | 1452

Example: Extended DHCPv6 Local Server Configuration | 1453

DHCPv6 Relay Agent | 1455

DHCPv6 Relay Agent Overview | 1455

DHCPv6 Relay Agent Options | 1456

Configuring DHCPv6 Relay Agent Options | 1457

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | 1458

Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets | 1460

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | 1462

DHCPv6 Monitoring and Management | 1464

Requesting DHCPv6 Local Server to Initiate Reconfiguration of Client Bindings | 1465

Verifying and Managing DHCPv6 Local Server Configuration | 1466

Verifying and Managing DHCPv6 Relay Configuration | 1467

5

Packet Triggered Subscriber Services**Packet Triggered Subscriber Services | 1470**

IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services | 1470

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1470

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1472

6

Address-Assignment Pools for Subscriber Management**Address-Assignment Pools for Subscriber Management | 1476**

Address-Assignment Pools for Subscriber Management | 1476

Address-Assignment Pools Overview | 1477

Address Allocation from Linked Address Pools | 1479

Address-Assignment Pool Configuration Overview | 1484

Configuring an Address-Assignment Pool Name and Addresses | 1486

Configuring a Named Address Range for Dynamic Address Assignment | 1486

Preventing Addresses from Being Allocated from an Address Pool | 1487

Configuring Address-Assignment Pool Usage Threshold Traps | 1489

Configuring Address-Assignment Pool Linking | 1490

Configuring Address-Assignment Pool Hold-Down | 1492

Configuring DHCP Local Address Pool Rapid Drain | 1493

Configuring Static Address Assignment | 1495

Configuring Duplicate IPv4 Address Protection for AAA | 1496

Example: Configuring an Address-Assignment Pool | 1498

Requirements | 1498

Overview | 1499

Configuration | 1499

7

DNS Addresses for Subscriber Management**DNS Addresses for Subscriber Management | 1503**

DNS Name Server Addresses for Subscriber Management | 1503

DNS Name Server Address Overview | 1503

Configuring DNS Name Server Addresses for Subscriber Management | 1505

Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | 1507

DNS Resolver for IPv6 DNS Overview | 1508

Configuring a DNS Server Address for IPv6 Hosts | 1508

M:N Subscriber Redundancy

M:N Subscriber Redundancy | 1511

M:N Subscriber Redundancy on BNG | 1511

M:N Subscriber Redundancy on BNG Overview | 1511

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization | 1543

Configure Subscriber Group Redundancy | 1544

Configure VRRP to Support M:N Redundancy | 1546

Configure Active Leasequery with Topology Discovery | 1547

How to Configure M:N Subscriber Redundancy with Pseudowires and DHCP Binding Synchronization | 1549

Configure Subscriber Group Redundancy | 1550

Configure Active Leasequery with Topology Discovery | 1553

Verifying M:N Redundancy and Active Leasequery Topology Discovery Information | 1554

M:N Subscriber Service Redundancy on DHCP Server | 1557

M:N Subscriber Service Redundancy on DHCP Server Overview | 1557

N+1 Support for BNG M:N Subscriber Service Redundancy | 1561

N+1 Support for BNG M:N Subscriber Service Redundancy Overview | 1562

How N+1 Support for BNG M:N Subscriber Service Redundancy Works | 1562

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | 1565

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Overview | 1565

How BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Works | 1566

Configuring BNG Redundancy Using Packet Triggered Recovery for DHCP Subscribers | 1568

Overview | 1568

Requirements | 1568

Topology | 1569

Configuration | 1570

Verification | 1614

Access Node Control Protocol and the ANCP Agent for Subscriber Services

Access Node Control Protocol and the ANCP Agent for Subscriber Services | 1621

ANCP Agent Neighbors and Operations | 1621

ANCP and the ANCP Agent Overview | 1622

ANCP Operations in Different Network Configurations | **1632**

Configuring the ANCP Agent | **1642**

Configuring ANCP Neighbors | **1644**

Associating an Access Node with Subscribers for ANCP Agent Operations | **1645**

Specifying the Interval Between ANCP Adjacency Messages | **1646**

Specifying the Maximum Number of Discovery Table Entries | **1647**

Configuring the ANCP Agent for Backward Compatibility | **1647**

Specifying How Long Processes Wait for the ANCP Agent Restart to Complete | **1648**

Configuring the ANCP Agent to Learn ANCP Partition IDs | **1649**

Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet | **1650**

Requirements | **1650**

Overview | **1651**

Configuration | **1658**

Verification | **1676**

ANCP Agent Traffic Shaping and CoS | **1681**

Traffic Rate Reporting and Adjustment by the ANCP Agent | **1682**

Preservation of CoS Shaping Across ANCP Agent Restarts | **1687**

Configuring the ANCP Agent to Report Traffic Rates to CoS | **1688**

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | **1693**

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | **1695**

Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | **1697**

Verifying and Monitoring CoS for ANCP Subscribers | **1699**

ANCP Agent and AAA | **1700**

ANCP Agent Interactions with AAA | **1701**

ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | **1703**

Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | **1713**

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | **1714**

ANCP Monitoring and Management | **1715**

Triggering ANCP OAM to Test the Local Loop | **1715**

Verifying and Monitoring ANCP Neighbors | **1717**

Clearing ANCP Neighbors | **1718**

Verifying and Monitoring ANCP Subscribers | 1719

Clearing ANCP Subscribers | 1720

Clearing and Verifying ANCP Statistics | 1721

Tracing ANCP Events for Troubleshooting | 1722

Configuring the ANCP Trace Log Filename | 1723

Configuring the Number and Size of ANCP Log Files | 1723

Configuring Access to the ANCP Log File | 1724

Configuring a Regular Expression for ANCP Messages to Be Logged | 1724

Configuring the ANCP Tracing Flags | 1725

Configuring the Severity Level to Filter Which ANCP Messages Are Logged | 1725

Diameter Base Protocol and its Applications

Diameter Base Protocol and its Applications | 1727

Diameter Base Protocol | 1727

Diameter Base Protocol Overview | 1728

Messages Used by Diameter Applications | 1731

Diameter AVPs and Diameter Applications | 1739

Configuring Diameter | 1762

Configuring the Origin Attributes of the Diameter Instance | 1763

Configuring Diameter Peers | 1763

Configuring the Diameter Transport | 1765

Configuring Diameter Network Elements | 1766

Example: Configure S6a Application | 1768

Requirements | 1768

Overview | 1768

Configuration | 1769

Verification | 1778

Gx-Plus for Provisioning Subscribers | 1781

Gx-Plus for Provisioning Subscribers Overview | 1782

Understanding Gx-Plus Interactions Between the Router and the PCRF | 1784

Configuring Gx-Plus | 1793

Configuring the Gx-Plus Partition | 1794

Configuring Gx-Plus Global Attributes | 1795

Provisioning Subscribers with Gx-Plus | 1796

Disabling PCRF Control of a Subscriber Session | 1797

3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1799

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1799
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1802
- Understanding Gx Interactions Between the Router and the PCRF | 1807
- Understanding Gy Interactions Between the Router and the OCS | 1819
- Gy File Backup Overview | 1826
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1828
- Understanding Upstream and Downstream Messages for the PCRF | 1833
- Configuring the OCS Partition | 1837
- Configuring the PCRF Partition | 1844
- Configuring OCS Global Parameters | 1851

NASREQ for Authentication and Authorization | 1852

- Diameter Network Access Server Application (NASREQ) | 1852
- Configuring the Diameter Network Access Server Application (NASREQ) | 1854

JSRC for Subscriber Provisioning and Accounting | 1856

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview | 1856
- Understanding JSRC-SAE Interactions | 1858
- JSRC Provisioning for Dual-Stack Subscribers | 1861
- JSRC Configuration Overview | 1865
- Configuring the JSRC Partition | 1865
- Assigning a Partition to JSRC | 1866
- Authorizing Subscribers with JSRC | 1867
- Provisioning Subscribers with JSRC | 1867
- Configuring JSRC for Dual-Stack Subscribers | 1867
- Excluding AVPs from Diameter Messages for JSRC | 1868
- Service Accounting with JSRC | 1869
- Configuring Service Accounting with JSRC | 1870

JSRC and Subscribers on Static Interfaces | 1872

- Subscribers on Static Interfaces Overview | 1872
- Subscribers over Static Interfaces Configuration Overview | 1876
- Example: Configuring Static Subscribers for Subscriber Access | 1877
- Specifying the Static Subscriber Global Access Profile | 1879
- Specifying the Static Subscriber Global Dynamic Profile | 1879
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers | 1880

- Configuring the Static Subscriber Global Authentication Password | **1880**
- Configuring the Static Subscriber Global Username | **1881**
- Creating a Static Subscriber Group | **1883**
- Specifying the Static Subscriber Group Access Profile | **1884**
- Specifying the Static Subscriber Group Dynamic Profile | **1884**
- Specifying the Static Subscriber Group Service Profile | **1884**
- Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group | **1885**
- Configuring the Static Subscriber Group Authentication Password | **1885**
- Configuring the Static Subscriber Group Username | **1886**

Monitoring and Management Diameter Information | **1888**

- Verifying Diameter Node, Instance, and Route Information | **1889**
- Verifying and Managing Diameter Application Information | **1890**
- Verifying and Managing Diameter Peer Information | **1892**
- Verifying Diameter Network Element Information | **1894**

Tracing Diameter Base Protocol Events for Troubleshooting | **1895**

- Configuring the Diameter Base Protocol Trace Log Filename | **1896**
- Configuring the Number and Size of Diameter Base Protocol Log Files | **1896**
- Configuring Access to the Diameter Base Protocol Log File | **1897**
- Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged | **1897**
- Configuring the Diameter Base Protocol Tracing Flags | **1898**
- Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged | **1898**

Troubleshooting Diameter Networks | **1899**

- Troubleshooting Diameter Network Configuration | **1899**
- Troubleshooting Diameter Network Connectivity | **1900**

Monitoring and Managing Static Subscriber Information | **1901**

- Forcing a Static Subscriber to Be Logged Out | **1901**
- Resetting the State of an Interface for Static Subscriber Login | **1901**
- Forcing a Group of Static Subscribers to Be Logged Out | **1902**
- Resetting the State of an Interface Group for Static Subscriber Login | **1902**
- Verifying Information about Subscriber Sessions on Static Interfaces | **1902**

Tracing Static Subscriber Events for Troubleshooting | **1903**

- Configuring the Static Subscribers Trace Log Filename | **1904**
- Configuring the Number and Size of Static Subscribers Log Files | **1904**

Configuring Access to the Static Subscribers Log File | 1905

Configuring a Regular Expression for Static Subscriber Messages to Be Logged | 1905

Configuring the Static Subscribers Tracing Flags | 1906

Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged | 1906

Configuration Statements and Operational Commands

dynamic-profile (Domain Map) | 1908

dynamic-profile (Static Subscribers) | 1909

Junos CLI Reference Overview | 1912

About This Guide

Use this guide to learn many aspects of configuring and connecting subscriber sessions, including the Junos OS AAA framework; using RADIUS or Diameter for authentication, service authorization, and accounting; CLI-based service activation/deactivation; DHCP and DHCPv6 for address assignment and client configuration; dual-stack access models; and managing subscriber access lines with ANCP.

1

PART

AAA for Subscriber Management

- AAA for Subscriber Management | 2
 - RADIUS for Subscriber Management | 646
 - Domain Maps for Subscriber Management | 825
 - Testing and Troubleshooting AAA | 846
 - RADIUS Dictionary Files | 859
-

AAA for Subscriber Management

IN THIS CHAPTER

- AAA Service Framework Overview | 2
- Standard and Vendor-Specific RADIUS Attributes | 3
- Subscriber Management RADIUS Dictionary Files | 645

AAA Service Framework Overview

The authentication, authorization, and accounting (AAA) Service Framework provides a single point of contact for all the authentication, authorization, accounting, address assignment, and dynamic request services that the router supports for network access. The framework supports authentication and authorization through external servers, such as RADIUS. The framework also supports accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.

When interacting with external back-end RADIUS servers, the AAA Service Framework supports standard RADIUS attributes and Juniper Networks vendor specific attributes (VSAs). The AAA Service Framework also includes an integrated RADIUS client that is compatible with RADIUS servers that conform to RFC-2865, *Remote Authentication Dial In User Service (RADIUS)*, RFC-2866, *RADIUS Accounting*, and RFC-3576, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*, and which can initiate requests.

You create the following types of configurations to manage subscriber access.

- Authentication—Authentication parameters defined in the access profile determine the authentication component of the AAA processing. For example, subscribers can be authenticated using an external authentication service such as RADIUS.
- Accounting—Accounting parameters in the access profile specify the accounting part of the AAA processing. For example, the parameters determine how the router collects and uses subscriber statistics. You can also configure AAA to enable the router to collect statistics on a per-service session basis for subscribers.

- **RADIUS-initiated dynamic requests**—A list of authentication server IP addresses in the access profile specify the RADIUS servers that can initiate dynamic requests to the router. Dynamic requests include CoA requests, which specify VSA modifications and service changes, and disconnect requests, which terminate subscriber sessions. The list of authentication servers also provide RADIUS-based dynamic service activation and deactivation during subscriber login.
- **Address assignment**—The AAA Service Framework assigns addresses to subscribers based on the configuration of local address-assignment pools. For example, the AAA framework collaborates with RADIUS servers to assign addresses from the specified pools.
- **Subscriber secure policy**—RADIUS VSAs and attributes provide RADIUS-initiated traffic mirroring on a per-subscriber basis.

RELATED DOCUMENTATION

[Configuring Access Profile Options for Interactions with RADIUS Servers | 653](#)

[Standard and Vendor-Specific RADIUS Attributes | 3](#)

[RADIUS Authentication and Accounting Basic Configuration | 721](#)

[RADIUS Accounting for Subscriber Access | 741](#)

Subscriber Secure Policy Overview

Standard and Vendor-Specific RADIUS Attributes

IN THIS SECTION

- [RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework | 4](#)
- [RADIUS IETF Attributes Supported by the AAA Service Framework | 4](#)
- [Juniper Networks VSAs Supported by the AAA Service Framework | 19](#)
- [AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs | 599](#)
- [AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs | 614](#)
- [DSL Forum Vendor-Specific Attributes | 626](#)
- [DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS | 637](#)
- [RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses | 642](#)
- [Support for Cisco Systems VSAs | 643](#)
- [Subscriber Management RADIUS Dictionary Files | 643](#)

- [Interface Text Descriptions for Inclusion in RADIUS Attributes | 643](#)

RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework

IN THIS SECTION

- [Benefits of Using RADIUS Standard Attributes and VSAs | 4](#)

The AAA Service Framework supports RADIUS attributes and vendor-specific attributes (VSAs). This support provides tunable parameters that the subscriber access management feature uses when creating subscribers and services.

RADIUS attributes are carried as part of standard RADIUS request and reply messages. The subscriber management access feature uses the RADIUS attributes to exchange specific authentication, authorization, and accounting information. VSAs allow the subscriber access management feature to pass implementation-specific information that provide extended capabilities, such as service activation or deactivation, and enabling and disabling filters.

When you use dynamic profiles, the AAA Service Framework supports the use of Junos OS predefined variables to specify the RADIUS attribute or VSA for the information obtained from the RADIUS server.

Benefits of Using RADIUS Standard Attributes and VSAs

- RADIUS standard attributes are necessary to communicate with an external RADIUS server for subscriber authentication, authorization, and accounting.
- Vendor-specific attributes extend the functionality of the RADIUS server beyond that provided by the public standard attributes, enabling the implementation of many useful features necessary for subscriber management and service support.

RADIUS IETF Attributes Supported by the AAA Service Framework

[Table 1 on page 5](#) describes the RADIUS IETF attributes that the Junos OS AAA Service Framework supports. Some attributes correspond to Juniper Networks predefined variables; see *predefined-variable-defaults (Dynamic Client Profiles)*



NOTE: A “Yes” entry in the Dynamic CoA Support column indicates that the attribute can be dynamically configured by Access-Accept messages and dynamically modified by CoA-Request messages. VSA's support Egress and Ingress policies for CoA.

Table 1: Supported RADIUS IETF Attributes

Attribute Number	Attribute Name	Description	Dynamic CoA Support
1	User-Name	<ul style="list-style-type: none"> Name of user to be authenticated. Configurable username override. Non-standard use for LLID preauthentication feature. 	No
2	User-Password	<ul style="list-style-type: none"> Password of user to be authenticated by Password Authentication Protocol (PAP). Configurable password override. Non-standard use for LLID preauthentication feature. 	No
3	CHAP-Password	<p>Value provided by a PPP (CHAP) user in response to the challenge.</p> <p>You can configure an override of the CHAP challenge response. When you configure an override CHAP password, the User-Password attribute contains the override, and the CHAP-Password attribute is not included in the Access-Request.</p>	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
4	NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.	No
5	NAS-Port	Physical port number of the NAS that is authenticating the user. For a tunneled PPP user in an L2TP LNS session, there is no physical port. In this case, the port value is reported as 4194303.	No
6	Service-Type	Type of service the user has requested or the type of service to be provided.	No
7	Framed-Protocol	Framing type used for framed access.	No
8	Framed-IP-Address	<ul style="list-style-type: none"> IP address to be configured for the user. 0.0.0.0 or absence is interpreted as 255.255.255.254. 	No
9	Framed-IP-Netmask	<ul style="list-style-type: none"> IP network to be configured for the user when the user is a router or switch to a network. Absence implies 255.255.255.255. 	No

Table 1: Supported RADIUS IETF Attributes (*Continued*)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
11	Filter-Id	<p>Name of a subscriber firewall filter, formatted as follows:</p> <ul style="list-style-type: none"> For an IPv4 input filter—IPv4-<i>ingress:ingress-filter-name</i> For an IPv4 output filter—IPv4-<i>egress:egress-filter-name</i> For an IPv6 input filter—IPv6-<i>ingress:ingress-filter-name</i> For an IPv6 output filter—IPv6-<i>egress:egress-filter-name</i> <p>RADIUS accounting request messages, Acct-Start and Acct-Stop, can include more than one Filter-Id attribute, one of each of the listed types.</p> <p>However, RADIUS Access-Accept messages can include only one attribute instance. The value is always treated as an IPv4 input filter name.</p>	Yes
12	Framed-MTU	Maximum Transmission Unit configured for the user, when it is not negotiated by some other means (such as PPP).	No
18	Reply-Message	<ul style="list-style-type: none"> Text that may be displayed to the user. Only the first instance of this attribute is used. 	No

Table 1: Supported RADIUS IETF Attributes *(Continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
22	Framed-Route	<p>String that provides routing information to be configured for the user on the NAS in the format:</p> <pre><addr>[/<maskLen>] [<nexthop> [<cost>]] [tag <tagValue>] [distance <distValue>]</pre> <p>If authd detects the IP address in the Framed-Route to be bad—for example, if the format is incorrect—the subscriber is not allowed to log in. Starting in Junos OS Release 19.1, the subscriber is allowed to log in, but without that route or the default route. For customers that use multiple framed routes, this behavior enables the subscriber to have partial access to the network using the routes that are accepted rather than not being allowed any access.</p> <p>Starting in Junos OS Release 18.2R1, if this attribute does not include the subnet mask, the MX Series router ignores the attribute but connects the session.</p>	No
24	State	String enabling state information to be maintained between the device and the RADIUS server.	No
25	Class	Arbitrary value that the NAS includes in all accounting packets for the user if supplied by the RADIUS server.	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
27	Session-Timeout	Maximum number of consecutive seconds of service to be provided to the user before termination of the session.	Yes Not supported for DHCP sessions.
28	Idle-Timeout	Maximum number of consecutive seconds of idle connection allowed to the user before termination of the session or prompt.	No
31	Calling-Station-ID	Phone number from which the call originated.	No
32	NAS-Identifier	NAS originating the request.	No
40	Acct-Status-Type	Whether this Accounting-Request marks the beginning of the user service (Start), the end (Stop), or the interim (Interim-Update).	No
41	Acct-Delay-Time	Number of seconds the client has been trying to send a particular record.	No
42	Acct-Input-Octets	Number of octets that have been received from the port during the time this service has been provided.	No
43	Acct-Output-Octets	Number of octets that have been sent to the port during the time this service has been provided.	No

Table 1: Supported RADIUS IETF Attributes *(Continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
44	Acct-Session-ID	<p>Unique accounting identifier that makes it easy to match start and stop records in a log file. The identifier can be in one of the following formats:</p> <ul style="list-style-type: none"> decimal—For example, 435264 description—In the generic format, <i>jnpr interface-specifier:subscriber-session-id</i>. For example, <i>jnpr fastEthernet 3/2.6:1010101010101</i> 	No
45	Acct-Authentic	Method by which user was authentication: whether by RADIUS, the NAS itself, or another remote authentication protocol.	No
46	Acct-Session-Time	Number of seconds that the user has received service	No
47	Acct-Input-Packets	Number of packets that have been received from the port during the time this service has been provided to a framed user.	No
48	Acct-Output-Packets	Number of packets that have been sent to the port in the course of delivering this service to a framed user.	No

Table 1: Supported RADIUS IETF Attributes *(Continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
49	Acct-Terminate-Cause	<p>Reason the service (a PPP session) was terminated. The service can be terminated for the following reasons:</p> <ul style="list-style-type: none"> • User Request (1)—User initiated the disconnect (log out). • Idle Timeout (4)—Idle timer has expired. • Session Timeout (5)—Client reached the maximum continuous time allowed on the service or session. • Admin Reset (6)—System administrator terminated the session. • Port Error (8)—PVC failed; no hardware or no interface. • NAS Error (9)—Negotiation failures, connection failures, or address lease expiration. • NAS Request (10)—PPP challenge timeout, PPP request timeout, tunnel establishment failure, PPP bundle failure, IP address lease expiration, PPP keep-alive failure, tunnel disconnect, or an unaccounted-for error. 	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
52	Acct-Input-Gigawords	Number of times the Acct-Input-Octets counter has wrapped around 2^{32} during the time this service has been provided. Can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update.	No
53	Acct-Output-Gigawords	Number of times the Acct-Output-Octets counter has wrapped around 2^{32} in the course of delivering this service. Can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update.	No
55	Event-Timestamp	Time that this event occurred on the NAS, in seconds, since January 1, 1970 00:00 UTC.	No
61	NAS-Port-Type	Type of physical port the NAS is using to authenticate the user. For a tunneled PPP user in an L2TP LNS session, there is no physical port. In this case, the port type is Virtual.	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
64	Tunnel-Type	<ul style="list-style-type: none"> Tunneling protocol to use (in the case of a tunnel initiator) or the tunneling protocol already in use (in the case of a tunnel terminator). Only L2TP tunnels are currently supported. 	No
65	Tunnel-Medium-Type	<ul style="list-style-type: none"> Transport medium to use when creating a tunnel for protocols that can operate over multiple transports. Only IPv4 is currently supported. 	No
66	Tunnel-Client-Endpoint	Address of the initiator end of the tunnel (LAC).	No
67	Tunnel-Server-Endpoint	Address of the server end of the tunnel (LNS).	No
68	Acct-Tunnel-Connection	Identifier assigned to the tunnel session. Value is the same as the Call Serial Number AVP received from the LAC in the ICRQ message.	No
69	Tunnel-Password	Encrypted password used to authenticate to a remote server. Recommended over using VSA Tunnel-Password [26-9] because of the encryption. Do not use both this attribute and the VSA.	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
77	Connect-Info	<ul style="list-style-type: none"> Information sent from the NAS that describes the subscriber's connection, such as transmit speed. Non-standard use for LLID preauthentication feature. 	No
82	Tunnel-Assignment -Id	Tunnel to which a session is assigned. When user profiles share the same values for Tunnel-Assignment-Id, Tunnel-Server-Endpoint, and Tunnel-Type, the LAC can group these users into the same tunnel. This grouping enables fewer tunnels to be created. (LAC)	No
83	Tunnel-Preference	<ul style="list-style-type: none"> Included in each set of tunneling attributes to indicate the relative preference assigned to each tunnel when more than one set of tunneling attributes is returned by the RADIUS server to the tunnel initiator. Included in the Tunnel-Link-Start, the Tunnel-Link-Reject, and the Tunnel-Link-Stop packets (LAC only). 	No

Table 1: Supported RADIUS IETF Attributes *(Continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
85	Acct-Interim-Interval	<p>Number of seconds between each interim accounting update for this session.</p> <p>The router uses the following guidelines for interim accounting:</p> <ul style="list-style-type: none"> • Attribute value is within the acceptable range (from 600 through 86,400 seconds)—Accounting is updated at the specified interval. • Attribute value of 0—No RADIUS accounting is performed. • Attribute value is less than the minimum acceptable value—Accounting is updated at the minimum interval (600 seconds). • Attribute value is greater than the maximum acceptable value—Accounting is updated at the maximum interval (86,400 seconds). <p>NOTE: Values are rounded up to the next higher multiple of 10 minutes. For example, a setting of 900 seconds (15 minutes) is rounded up to 20 minutes (1200 seconds).</p>	No

Table 1: Supported RADIUS IETF Attributes *(Continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
87	NAS-Port-Id	<p>Text string that identifies the physical interface of the NAS that is authenticating the user.</p> <p>For a tunneled PPP user in an L2TP LNS session, there is no physical port, and the NAS-Port-Id value has the following format: <i>media:local address:peer address:local tunnel id:peer tunnel id:local session id:peer session id:call serial number</i>. For example, lp:198.51.100.1:192.168.0.2:3341:21031:16138:11846:2431. The local information refers to the LNS and the peer information refers to the LAC.</p>	No
88	Framed-Pool	Name of an assigned address pool to use to assign an address for the user.	No
90	Tunnel-Client-Auth-Id	Name of the tunnel initiator (LAC) used during the authentication phase of tunnel establishment.	No
91	Tunnel-Server-Auth-Id	Name of the tunnel terminator (LNS) used during the authentication phase of tunnel establishment.	No
95	NAS-IPv6-Address	Address of the NAS that is requesting authentication of the user.	No
96	Framed-Interface-ID	Interface identifier that is configured for the user.	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
97	Framed-IPv6-Prefix	IPv6 prefix and address that are configured for the user. Prefix lengths of 128 are associated with host addresses. Prefix lengths less than 128 are associated with NDRA prefixes.	No
98	Login-IPv6-Host	System the user connects to when the Login-Service attribute is included.	No
99	Framed-IPv6-Route	IPv6 routing information that is configured for the user.	Yes
100	Framed-IPv6-Pool	Name of the assigned pool used to assign the address and IPv6 prefix for the user.	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
101	Error-Cause	<p>Reason that the RADIUS server does not honor Disconnect-Request or CoA-Request messages. Depending on the value, can be included in CoA NAK or Disconnect NAK messages.</p> <ul style="list-style-type: none"> • 201—Residual Session Context Removed (Disconnect ACK only) • 202—Invalid EAP Packet (Ignored) • 401—Unsupported Attribute; request contains unsupported attribute. • 402—Missing Attribute; critical attribute missing from request • 403—NAS Identification Mismatch • 404—Invalid Request • 405—Unsupported Service • 406—Unsupported Extension • 407—Invalid Attribute Value • 501—Administratively Prohibited • 502—Request Not Routable (Proxy) • 503—Session Context Not Found 	No

Table 1: Supported RADIUS IETF Attributes (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
		<ul style="list-style-type: none"> • 504—Session Context Not Removable • 505—Other Proxy Processing Error • 506—Resources Unavailable • 507—Request Initiated • 508—Multiple Session Selection Unsupported 	
123	Delegated-IPv6-Prefix	IPv6 prefix that is delegated to the user.	No
168	Framed-IPv6-Address	IPv6 address of the authenticated user. The Framed-IPv6-Address attribute is sent if the IPv6 address is assigned to the subscriber.	No
242	Ascend-Data-Filter	Binary data that specifies RADIUS policy definitions.	Yes

Juniper Networks VSAs Supported by the AAA Service Framework

Table 2 on page 20 describes Juniper Networks VSAs supported by the Junos AAA Service Framework. The AAA Service Framework uses vendor ID 2636, which is assigned to Juniper Networks by the Internet Assigned Numbers Authority (IANA). Some VSAs correspond to Juniper Networks predefined variables; see *Predefined Variables That Correspond to RADIUS Attributes and VSAs*.



NOTE: A “Yes” entry in the Dynamic CoA Support column indicates that the attribute can be dynamically configured by Access-Accept messages and dynamically modified by CoA-Request messages.

Table 2: Supported Juniper Networks VSAs

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-1	Virtual-Router	Client configuration logical system name. Allowed	string: logical system name	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		only from AAA server for default logical class system:router output		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		g in st a n c e. W h e n t h i s V S A i s n o t i n c l u d e d i n t h e s u b s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		cr ib e r p r of il e, t h e r o ut in g in st a n c e a s si g n e d t o t h e s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Subscriber — the one in which the subscriber session		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		comes up – varies by subscription type. For DHCP		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		and PPPoE subscribers, it is the default routing instance		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		For L2TP tunnels, it is the routing instance.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ance in which the tunnel endpoint is, whether the default		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		non-default. If the tunnel endpoint is not defined		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		default and you want the L2TP session to be in the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		fault routing instance, you must use the Virtual-Router		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Server VSA Attributes to set the desired routing instance.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-4	Primary-DNS	Client identifier: 4-byte hexadecimal primary address assigned during IP C/P.	integer: 4-byte hexadecimal address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-5	Secondary-DNS	Client identifier: 4-byte hexadecimal string consisting of 16 hexadecimal characters		No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-6	Primary-WINS	Client identifier: WINS (NBNS) address	integer: 4-byte primary address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		P C P.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-7	Secondary-WINS	Client identifier: WINS (NetBIOS) address that is generated during login	4-byte hexadecimal	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		P C P.		
26-8	Tunnel-Virtual-Router	V i r t u a l r o u t e r n a m e f o r t u n n e l c o n n e c t i o n.	s t r i n g: <i>tu n n e l - vi rt u a l - ro ut er</i>	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-9	Tunnel-Password	Tunnel-Password is a text attribute. Don't use both	string: tunnel-password	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		his VSA and the standard RADIUS attribute Tunnel-		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Password [69]. We recommend that you use the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		standard attribute to be caused by the password is encryption		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		present when that attribute is used.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-10	Ingress-Policy-Name	Input policy name to apply to client interface.	string: input-policy-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-11	Egress-Policy-Name	Output group policy - policy name to apply to client in the area.	string: output-policy-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-12	Policer-L1-Input-Gigawords	(Junos OS Evolved) Number of times the policer-L1	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input-Output statistics counter		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable field while the policy server is active.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-13	Policer-L1-Input-Octets	(Junos OS Evolved) Number of octets received on t	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		helps the policy enforcement of the course of delivery of this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to a frame dissector.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-14	Policer-L2-Input-Gigawords	(Junos OS Evolved) Number of times the policer-L2	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input-Output statistics counter		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable field while the policy server L2 session is active.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-15	Policer-L2-Input-Octets	(Junos OS Evolved) Number of octets received on t	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		he L2 policer in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to a frame dissector.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-16	Policer-L3-Input-Gigawords	(Junos OS Evolved) Number of times the policer-L3	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input-Output statistics counter		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Identifies the policy server L3 service is active.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-17	Policer-L3-Input-Octets	(Junos OS Evolved) Number of octets received on t	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		he L3 policer in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to a frame dissector.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-18	Policer-L4-Input-Gigawords	(Junos OS Evolved) Number of times the policer-L4	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input-Output statistics counter		
		etherhash-wrap-pendarray-round4-oct		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable field while the policy server L4 session is active.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-19	Policer-L4-Input-Octets	(Junos OS Evolved) Number of octets received on t	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		he L4 policer in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to a frame dissector.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-20	Policer-L1-Input-Gigapackets	(Junos OS Evolved) Number of times the policer-L1	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input packets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-21	Policer-L1-Input-Packets	(Junos OS Evolved) Number of packets received	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		not the L1 policy enforcement of the course of delivery of this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		create a framed user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-23	IGMP-Enable	Whether the router IGMP is enabled or disabled on a client interface	Integer: <ul style="list-style-type: none">• 0 = disabled• 1 = enabled	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		effective.		
26-24	PPPoE-Description	Client Mapping Address	string: popoecli-nt-macs-address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-25	Redirect-VRouter-Name	Client's logical address system's remote router outgoing	string: logical address system's router outgoing	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to which logical systems mirror outgoing instances and their request		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		redirect for user authentication.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-29	Queue-0-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue-0	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		output packet size attribute rollover		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		el d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-30	Tunnel-Nas-Port-Method	Method used to determine whether the RADIUS	4-octet integer: <ul style="list-style-type: none">• 0 = none• 1 = Cisco COLD	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		server conversion system to the LLNS that the physical NASS port		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		number of physical ports		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		, such has Ethernet network or ATM. This information is con		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		viewed only when the VSA value is 1. The VSA is for m		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		attends such that the first octet indicates the tunnel		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		element that remains in the network for the entire session		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		v a l u e.		
26-31	Service-Bundle	S S C s e r v i c e b u n d l e.	s t r i n g <i>b u n d l e</i> <i>e- n a m e</i>	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-32	Queue-0-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue out of the course delivery system to access		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-33	Tunnel-Max-Sessions	Maximum number of sessions allowed in a tunnel.	Integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-34	Framed-IP-Route-Tag. Supported only on JunosE for ERX and E320 platforms.	Router tag: 4-octet or less returned from the ip address	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		r e s s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-42	Input-Gigapackets	Number of times the input-packets at tributaries	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		version 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-43	Output-Gigapackets	Number of times the output packets at attribute rol	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Is overrits 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-47	Ipv6-Primary-DNS	Client's primary IPv6 DNS address	hexadecimal string: <i>ip v 6- primary address</i>	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		D H C P.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-48	Ipv6-Secondary-DNS	Client's IPv6 Secondary DNS Address	hexadecimal IPv6 address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dynamic DHCP.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-49	Queue-1-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue-1	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		output packet size attribute rollover		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		el d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-50	Queue-1-Output-Packets	(Junos OS Evoolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue 1 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-51	Disconnect-Cause	Description of the disconnect cause	hexadecimal value: <i>disconnect-cause</i>	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		connected, and L2TP layer of the LNS initiates the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		etermine minimum attribute. The PPP Disconnection Cause Code (L2TP		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		AVP 46) is included in VSA 26-51 in the Accounting - S		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to process a message that the router sends to the RADIUS		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		IUS server.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-55	DHCP-Options	Client Default DHCP Option s. Setting in Junos OS Release	hexadecimal decimal string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		17.4 R1, includes only DHCPv4 options. In a reliable manner		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		es, including sbooth DHCP Cpv4 and DHCP Cpv6 options.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-56	DHCP-MAC-Address	Client MAC Address	string: mac-address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-57	DHCP-GI-Address	DHCP-Relay Agent IP address.	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-58	LI-Action	Tr aff ic m ir r o ri n g a ct io n. F o r d y n a m ic C o A , V S A 2 6 - 5	sa lt- e n cr y pt e d in te g er 0 = st o p m irr or in g 1 = st ar t m irr or in g	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8c h a n g e s t h e a c t i o n o n t h e m i r r o r e d t r a f f i c i d e n t i f i e d	2 = n o a c t i o n	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		by VSA 26-59. CoA-Request message that includes		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dean y o f t h e R A D I U S - b a s e d m i r r o r i n g a t t r i b u t e s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		(VSAs 26-58, 266-59, 266-60, or 266-61) must always		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		includes all our VSA's. If the CoA action is to stop mir		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		originating (VSA 26-58 value is 0), then the value of the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Other relevant attributes in the CoA message must		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		attach the existing attribute values, or the action fails.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-59	Med-Dev-Handle	Identifies the cryptographic header string associated with mirrored or redirected traffic to a specific	salt-encrypt-header string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration. For dynamic CoA, VSA 26-58 changes		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gest the action on the mirror or redirect traffic identified by VSA		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		26-59. CoA-Request message that includes any		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		of the RADUIS-based mirror originating attributes (VSAs)		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		26-58, 26-59, 26-60, or 26-61)		
		must always include		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		all four VSA s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-60	Med-Ip-Address	IP address record associated with the user's session	sa	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		chmirrored traffic is forwarded. CoA-Request messages		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sagest that include any of the RADIOS-based m		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		routing attributes (VSAs 26-58, 66-92, 60, or 26		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- 61) must always include all four VSA s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-61	Med-Port-Number	UDP port in the header of the connection to the destination in a notation of device to which	sa-lt-pn-crypt-enabled-integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		h m i r r o r e d t r a f f i c i s f o r w a r d e d. C o A - R e q u e s t m e s s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		aggest that include any of the RADIOS-based mirror		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		orig attributes (VSAs 266-58, 266-59, 266-60, or 266-		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		61) must always include all four VSAs.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-63	Interface-Desc	Text string: interface description	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sin t e r f a c e.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-64	Tunnel-Group	Name of the tunnel group (profile) assigned to the user	string: tunnel-group-name	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		o m a i n m a p.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-65	Activate-Service	Service: service-name activate for the subscriber. Tagged V	string: service-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		SA, which supports 8 tags (1-8).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-66	Deactivate-Service	Service description for the service to be deactivated.	string: service-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		server is deactivated when the volume is exceeded	limit	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		d. Tagged VSA, which supports 8 tags (1-8).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-68	Service-Timeout	Number of seconds condition should last at the server side and act	integer <ul style="list-style-type: none">range = 0 to 1677215 seconds	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ive; se rvi c e is d e a ct iv a t e d w h e n t h e t i m e o u t e x p i r e	• 0 = n o t i m e o u t	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		s. Tagged VSA, which supports 8 tags (1-8).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-69	Service-Statistics	Whether the router statistics for the service are visible or not	<ul style="list-style-type: none">0 = disabled1 = enabled	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		stable. Tagged VSA, which supports 8 tags (1-	• 2 = enableable time and volume statistics	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8).		
26-71	IGMP-Access-Name	Access list to use for the group (G) filter.	string: 32-bit	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-72	IGMP-Access-Src-Name	Access source list of users for the session source - group (S, G) file	string: 32-octet	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ter.		
26-74	MLD-Access-Name	Access list to use for the group (G) filter.	string: 32-bit octet	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-75	MLD-Access-Src-Name	Access source list of users for the session - group (S, G) file	string: 32-octet	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		t e r.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-77	MLD-Version	MLD-Deprioritization. ol version.	MLD-Deprioritization. ol version. • 1 = MLD version 1 • 2 = MLD version 2	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			on	2

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-78	IGMP-Version	IGMP protocol collection version.	Integer: 1-255 <ul style="list-style-type: none">1 = IGMP version 12 = IGMP version 2	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-83	Service-Session	Name of the service.	string: service-name	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-84	Policer-L2-Input-Gigapackets	(Junos OS Evolved) Number of times the policer-L2	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input packets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-85	Policer-L2-Input-Packets	(Junos OS Evolved) Number of packets received	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		not the L2 Policy in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		create a framed user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-86	Policer-L3-Input-Gigapackets	(Junos OS Evolved) Number of times the policer-L3	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input packets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-87	Policer-L3-Input-Packets	(Junos OS Evolved) Number of packets received	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		not the L3 policy in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		create a framed user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-88	Policer-L4-Input-Gigapackets	(Junos OS Evolved) Number of times the policer-L4-3	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- Input packets attributed to router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-89	Policer-L4-Input-Packets	(Junos OS Evolved) Number of packets received	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		not the L4 policy in the course of delivering this service		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		create a framed user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-91	Tunnel-Switch-Profile	Tunnelling: description of the tunnel switch profile	string: profile name	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		assubscribere session switch checked to assess conditions		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		on the remote LNS. Takes precedence over return		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		network switch chip processor of the switch applied in any other manner.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-92	L2C-Up-Stream-Data	Actual upstream rate of traffic (ASCI) (A	string: actual upstream rate of traffic (ASCI)	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		SCILL encoding) as defined in GSSMP extension for L	o de d)	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Layer 2 control (L2C) topology discovery and Line Co		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Configuration		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-93	L2C-Down-Stream-Data	Actual download speed in bits per second (bps)	string: actual download speed in bits per second (bps)	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		(Access Control List entry coded as defined in GSMP Extension for	n	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Layer 2 control (L2C) topology discovery and Line		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Configuration.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-94	Tunnel-Tx-Speed-Method	Method to determine the source of the traffic	Integer: 4-octet • 0 = none • 1 = static L2 • 2 = dynamic	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		he transactional system limits peer-to-peer discovery of devices. Overriding global configuration	dynamic transactional system method is not supported	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ration in the CLI.	ted; these static Layer 2 methods do not use dynamic	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-97	IGMP-Immediate-Leave	IGMP: Immediate Leave. Attribute value.	<div><div>in</div><div>ter</div><div>ger</div><div>:</div><div>4-</div><div>o</div><div>ct</div><div>et</div><div>• 0</div><div>=</div><div>d</div><div>i</div><div>s</div><div>a</div><div>b</div><div>l</div><div>e</div><div>• 1</div><div>=</div><div>e</div><div>n</div><div>a</div><div>b</div><div>l</div><div>e</div></div>	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-100	MLD-Immediate-Leave	MLD-Immediate-Leave: 4-octet value.	<div><div>in</div><div>ter</div><div>:</div><div>4-</div><div>octet</div><div>• 0 = disabled</div><div>• 1 = enable</div></div>	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-106	IPv6-Ingress-Policy-Name	Input policy name to apply to a user's IPv6 interface	string: policy-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		c e.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-107	IPv6-Egress-Policy-Name	Output policy name to apply to a user's IPv6 interface	string: policy-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		a c c e.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-108	CoS-Parameter-Type	CoS traffic affects - , sh a pi n g p a r a m e t e r t y p e a n d d e s c r i p t i o n:	Tr o p a r t s , d e l i m i t e d b y w h i t e s p a c e: <ul style="list-style-type: none">• P a r a m e t e r t	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		<ul style="list-style-type: none">• T01: Shared user limit - maximum number of sessions• T02: Shared user limit - maximum number of sessions	By default, the value is 1. The value can be set to 1 or 2.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			<ul style="list-style-type: none">• T06: Traffic - global• T05: Exchange - remote mode• T06: Traffic	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- control plane • T07: Shaping mode • T0		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			8 : Byte add just	
		• T09 : Add just minimum		
		• T10		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			: E x c e s s - r a t e h i g h	
		• T 1 1 : E x c e s s - r a t e		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			low	
		<ul style="list-style-type: none">• T12: Shaping rate bursts• T13: Guarantee		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			not needed rate bursts	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-109	DHCP-Guided-Relay-Server	Indicates the IP address of the DHCP server that the DHCP relay agent	integer: 4-byte IP address	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		element uses test of forward the device cover PDU s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-110	Acc-Loop-Cir-Id	Identifier for the access loop circuit	string: up to 63 ASCII characters	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		other access nodes.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-111	Acc-Aggr-Cir-Id-Bin	Unique identifier for the DSL line.	integer : 8-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		in the following examples:		
		<ul style="list-style-type: none">• Ethernet access tag		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Registration — Ethernet — Settings / Port [:]		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		<ul style="list-style-type: none">ATM aggregation — atmsslot/port: vpi.vv		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
			<i>c i</i>	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-113	Act-Data-Rate-Up	Actual upstream rate of the session	Integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		notification		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-114	Act-Data-Rate-Dn	Actual download rate of the session's	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		synchronized DSS Link.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-115	Min-Data-Rate-Up	Minimum number of upstream packets received for the	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		subscriber.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-116	Min-Data-Rate-Dn	Minimum download rate in kbps	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hesubscribe		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-117	Att-Data-Rate-Up	Maximum upstream port rate that the subscriber	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rcan attain.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-118	Att-Data-Rate-Dn	Maximum download rate in bytes per second	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		iber can attain.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-119	Max-Data-Rate-Up	Maximum upstream port speed in megabits per second	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		subscriber.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-120	Max-Data-Rate-Dn	Maximum download rate in bits per second	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hesubscriber.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-121	Min-LP-Data-Rate-Up	Minimum number of upstream packets received at a rate in low power state	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-122	Min-LP-Data-Rate-Dn	Minimum download rate in low power state	Integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration for the subscriber.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-123	Max-Interlv-Delay-Up	Maximum number of ways up stream reachability	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-124	Act-Interlv-Delay-Up	Subscriber's actual online - way up stream area in the network	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		g d el a y. .		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-125	Max-Interlv-Delay-Dn	Maximum number of consecutive ways down stream element delay	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration for the subscriber.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-126	Act-Interlv-Delay-Dn	Subscriber's actual online - way down stream rate limit release	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		vi n g d el a y.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-127	DSL-Line-State	Statistics of the DSL line.	Integer: 4-octet <ul style="list-style-type: none">• 1 = Show up time• 2 = Idle• 3 = Sil	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			en	t

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-128	DSL-Type	Encryption algorithm used by the subscriber's computer	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		down with the DSLAM interface from which the request is sent		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		initiated.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-130	Qos-Set-Name	Interface access egress to application profile.	string: interface access egress to the dynamic profile.	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-131	Queue-2-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 2 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-132	Queue-2-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue 2 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-133	Queue-3-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 3 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packet classification attribute number of the 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-134	Queue-3-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue 3 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-135	Queue-4-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 4 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packet classification attribute number		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-136	Queue-4-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue 4 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-137	Queue-5-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 5 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packet counts attribute value rollover interval - octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-138	Queue-5-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue size in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-139	Queue-6-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 6 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packet keys attribute identifier. 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-140	Service-Interim-Acct-Interval	Amount of time between updates	<ul style="list-style-type: none">Range = 600 to 640 seconds0 = disabled	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		these for this session record. Tagged VSA, which supports	able to be used. NOTE: Value used as reference only.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8 tags (1-8).	xth highest hexadecimal value of 10 minutes. For example, 1	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			enable, allow, session, timing, group, configuration, 9000, 000, session, conditions, (15 minutes), utilities,) is required	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			nn dd ee dd uu pp tt cc 22 CC mm ii nn uu tt ee ss (11 22 CC CC ss ee cc cc nn dd ss) .	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-141	Downstream-Calculated-QoS-Rate	Calculated downstream rate (adjusted upstream rate) in Kbps	range = 1000 to 999967295	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		assess the ANCP configuration. A change in value		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		error results in a non-immediate data item in the queue		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		st .		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-142	Upstream-Calculated-QoS-Rate	Calculated upstream rate (adjusted upstream pipe rate in Kbps)	range = 1000 to 9999967295	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		set by the ANCP configuration. A change in value		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Results in a minimum immediate attribute in the Internet - Accounting in the require		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		st .		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-143	Max-Clients-Per-Interface	Maximum allowed wireless client sessions per interface. For	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rr DHCP Clients, this value is the maximum sessions per		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Logical interface. For PPPoE clients, this value is the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		maximum sessions (PPPoE interface access) per PPPoE user		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Indefinitely in the effective.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-145	Queue-6-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame user.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-146	CoS-Scheduler-Pmt-Type	CoS Scheduler parts, delimiter, item name, type, space, and description:	Throttle parts, delimiter, item name, type, space, and description: • Scheduler	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		<ul style="list-style-type: none">• N u l : C o S s c h e d u l e r n a m e• T O 1 : C o S s c h e d E x	<p>n a m e P a r a m e t e r t y p e P a r a m e t e r v a l u e</p>	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			ua m epl res t: r • a n s m i t r a t • e • T O 2 : C o S s c h e d u l e r b	b e - s c h e d T O 1 1 2 m b e - s c

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			uffe r s i z e	h e d T 0 2 2 6
		• T 0 3 : C o S s c h e d u l e r p r i o r i t y		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
		<ul style="list-style-type: none">• T04: CoS scheduled order drop - profile allow• T0	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			5 : CoS scheduler drop - profile medium -	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			low	
		<ul style="list-style-type: none">• T06: CoS schedule drop - profile		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			dynamic - high	
		<ul style="list-style-type: none">• T07: CoS scheduler drop - p		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	profile-high	
		•	T08: CoS scheduled order drop-pr	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			optional	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-149	Queue-7-Output-Gigapackets	(Junos OS Evolved) Number of times the Queue 7 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output packet classification attribute value rollover interval - octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-151	IPv6-Acct-Input-Octets	IPv6 received octets.	integer	No
26-152	IPv6-Acct-Output-Octets	IPv6 transmitted octets.	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-153	IPv6-Acct-Input-Packets	IPv6 received packets.	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-154	IPv6-Acct-Output-Packets	IPv6 traffic sent	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-155	IPv6-Acct-Input-Gigawords	IPv6 received gigawords.	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-156	IPv6-Acct-Output-Gigawords	IPv6 traffic transmitted in gigawords.	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-158	PPPoE-Padn	Routing table added for PPPoE sessions	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-159	L2tp-Csun-enable	(Junos OS Evolved) Overwrite the configuration for	Integer <ul style="list-style-type: none">0 = disabled1 = enabled	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		LNSE endpoints that have not supported L2TP Connection S		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		peer update Notification (CSUN) as defined in RFC 5		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		515.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-160	Vlan-Map-Id	Trunk VLAN tag corresponding to the core-facination	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gtrunk phy-sical interface. VLAN-Mapped Id (26-160),		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		In n e r- V la n - M a p - Id (2 6 - 1 8 4) , a n d C o r e - F a c i n g - In		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Interface (26-185) collectively represent the network		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		work service provider organization for the subs		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		describe for the Layer 2 cross-connect in a Layer 2 wh		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ole s al e c o nf ig u r a ti o n.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-161	IPv6-Delegated-Pool-Name	Address reserved for pool of locally allocated addresses	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		prefix (IP Address).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-162	Tx-Connect-Speed	Indication of transmission speed of the user's connection	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		io n.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-163	Rx-Connect-Speed	Indication of receive speed of the user's connection	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		io n.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-164	IPv4-Release-Control	Indicates the status of the server's release control	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		essential location and deallocation.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		it in g p a r a m e t e r t y p e a n d d e s c r i p t i o n:	<ul style="list-style-type: none">• P a r a m e t e r v a l u e• E x a m p l e s :<ul style="list-style-type: none">• T 0 1 : H i g• T 0 1 5 0 k	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		h• T c 0 o 6 m 1 m k i tt e d i n f o r m a t i o n r a t e . T h e c o m m i tt		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			defined information rate (CIR) for high-prior	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
		• i t y t r a ff c i n a h i e r a r c h i c a l p o l i c e r . TO	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			2: High max comm itted infor mation rate.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			The maximum number of sessions that can be established on the interface.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			high-priority traffic, which allows rat	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			eliminating support of this value.	
		• T03: High com		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			limited bursts size	
			The recommended	
			limited bursts size	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			efor high-priority traffic, which limits	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			the amount of burst traffic allowed.	
		• T04:		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Med - high committ ed information rate The		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			Configured	
			Priority	
			Traffic	
			• TO	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			5 : Med-high maximum committed information rate	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			the The maximum CIR formed medium -high priority	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			priority traffic	
			• TOS: Medium-high committed bandwidth	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			stsize The burst size of the formed medium-high-p	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Priority traffic.		
		• T07: Med-low committ		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
		Information rate The CIR for medium - low p	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Priority traffic.		
		• TOS: Medium-low maximum		
		mitt		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			edit information rate	
			The maximum CIR for	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Dynamic CoA Support
		ed i u m - l o w p r i o r i t y t r a f f i c . • T O 9 : M e d - l o	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			with community attribute defined by users to restrict size of the bursts to size of the bursts to size of the bursts	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			medium-low priority traffic.	
		•	T10: low co	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			mmited information rate. The CIR for low-	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			priority traffic.	
		• T11: low maximum permitted		
		id		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Information rate	The maximum CIR for low-	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			priority traffic	
			T12: low committeds	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			tsize The burst size for low-priority	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			raffc.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-173	Service-Activate-Type	Indicates the type of service activation. This is a tag used to indicate the type of service activation. The value is a tag used to indicate the type of service activation. The value is a tag used to indicate the type of service activation.	Integer: 4-octet • 1 = dynamic activation. The value is a tag used to indicate the type of service activation.	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		data attribute.	denial	
			• 2	
			=	
			o	
			p	
			-	
			s	
			c	
			r	
			i	
			p	
			t	
			f	
			o	
			r	
			b	
			u	
			s	
			i	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				nesses service vices

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-174	Client-Profile-Name	Enable RADIUS to override the assigned signing client identity	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		name of interface included in client profile of IEEE - n ame -		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		string. Enabled RADIUS to distinguish users differently		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		amic pr of il e s u s e d o n t h e r o u t e r w h e n t h e v e r s i o n		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		- alias- string is included. The version - alias- string is confirmed		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ig u r e d, al o n g w it h t h e d y n a m ic p r o f il e n a m e is p r o vi d		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		edit or RADILUS through a unique s-Account or Account		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ess-R equ st m e s s a g e.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-177	Cos-Shaping-Rate	Effective down stream rate for subscribers	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e r.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-178	Action-Reason	Indicates the reason for the action. The value is a string that represents the reason for the action. The value is a string that represents the reason for the action. The value is a string that represents the reason for the action.	String: "100 In progress" "104 Service active" "120 Service" "120 Service"	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Typical usage of the CoA failure message.	Not supported. "1" indicates the initial processing of the message.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			formatted" "1 2 3 N o s e r v i c e s " "1 2 4 S e r v i c e s l i m i t e d e d"	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			"125 Bulk request message size limit exceeded"	
			"128 Maximum	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			current CoAs"	
			"126 CoA request time out"	
			"127 Log	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			output progress"	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-179	Service-Volume-Gigawords	Amount of traffic, in 4 GB units, that has been consumed by the session.	Integer • range = 0 to 16,772,154 GB units	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		revise; server virtualization is deactivated when the volume is exc	<ul style="list-style-type: none">0 = no limit	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ended. Tagged VSA, which supports 8 tags (1-		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-180	Update-Service	New version available service notification and time quota status for existing	string: service-name	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		g s e r v i c e . T a g g e d V S A , w h i c h s u p p o r t s 8 t a g s (1 -		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8).		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-181	DHCPv6-Guided-Relay-Server	Indicates whether the DHCPv6 server is configured as a guided relay server. The value is <i>ip v6-addresses</i> if the DHCPv6 server is configured as a guided relay server. The value is <i>no</i> otherwise.	ip v6-addresses	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		CPv6 relay agent for forwarding packets to the Solicitor and subscriber		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		queue n t P D U s. U s e m ult ip le in st a n c e s o f t h e V S A t o s p e		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration of servers.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-182	Acc-Loop-Remote-Id	Reporting the ANCP Access-Loop-Remote-Id attribute	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Attribute.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-183	Acc-Loop-Encap	Represents the encapsulation of the loopback address in the Access Point Name (APN) configuration.	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tribute.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-184	Inner-Vlan-Map-Id	Inner VLAN tag allocated from the range	integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		one doesn't have core-facilitating phy-sical interface, used to		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		oswap (replace) the autoselected VLAN tag on the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		access in the reference. VLAN - Mapping Id (2601), Inner V		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		lan - Mapping - Id (26 - 184), and Core - Fabric - Networking - Interface		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		26-185) collectively represent the network works		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Service Provider Virtual Edge RFA Client Location Information for the subscriber		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Port the Layer 2 cross-connect in a Layer 2 whole sale		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configuration.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-185	Core-Facing-Interface	Number of the core-facing physical interface	string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		forwarding table Layer 2 whole session's down stream		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		request to the network		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Service Provider (NSP) Router. VLAN-Mapping Id (26-		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		160), In n e r- V la n - M a p - Id (2 6 - 1 8 4), a n d C o r e - F a ci		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ng-Interface (26-185) collectively represent the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable network work services provided for the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		subscribers for the Layer 2 cross-connect in a Layer		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		required whole scale configuration.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-189	DHCP-First-Relay-IPv4-Address	IP v4 address of the first relay agent of the client/serve	integer : 4-byte <i>ip - address</i>	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		review binding.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-190	DHCP-First-Relay-IPv6-Address	IP v 6 address of the first relay	hexadecimal string: ip v 6-addresses	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		review binding.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-191	Input-Interface-Filter	Notification of an input filter to be attached to a checked to a firewall	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		any interface.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-192	Output-Interface-Filter	Name of output filter to be attached to ached to a	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ilyan interface.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-193	Pim-Enable	Enable PIM on the interface. PIM is used for multicast routing. PIM is enabled on the interface by default.	Integer: 4-octet • 0 = disabled • 1 = enabled	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				enable

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-194	Bulk-CoA-Transaction-Id	Accounting transaction identifier for the session	integer : 4-octet	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sof related CoA Request status attribute. This attribute		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ute is untagged and value 0 is reserved.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-195	Bulk-CoA-Identifier	A unique identifier for each CoA Request message sent	integer : 4-octet	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		that is part of the same transaction as specified by t		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		he Bulk - CoA - Transaction - Id VSA . This attribute is un		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tagged and the value is reserved.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-196	IPv4-Input-Service-Set	Notification of a new IPv4 input service set to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-197	IPv4-Output-Service-Set	Naming of an IPv4 output service set to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		h e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-198	IPv4-Input-Service-Filter	Notification of a new IPv4 input service filter to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-199	IPv4-Output-Service-Filter	Notification of an IPv4 output service filter to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		c h e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-200	IPv6-Input-Service-Set	Notification of a new IPv6 input service received to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-201	IPv6-Output-Service-Set	Naming of an IPv6 output service set to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		h e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-202	IPv6-Input-Service-Filter	Name of an IPv6 input service filter to attach	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e d.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-203	IPv6-Output-Service-Filter	Naming of an IPv6 output service filter to be attached	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ch e d.		
26-204	Adv-Pcef-Profile-Name	Na m e o f a P C E F p r o f i l e t o b e a t t a c h e d.	st ri n g	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-205	Adv-Pcef-Rule-Name	Name of a PCC rule to activate.	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-206	Reauthentication-On-Renew	Reauthentication-On-Renew	Integer <ul style="list-style-type: none">0 = disabled1 = enabled	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ted.	ation when DDHC Prenew request is received	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			from the client network	
			• all other reserved values = invalid	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
			i d	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-207	DHCPv6-Options	Description of DHCPv6 options	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Whether RADIOS server uses reverse TLS V options. In release		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		available in Junos OS Release 17.1R1, this VSA is not		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Not supported. DHCPv6 option 82 is not included in the standard in 2		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		6 - 55, DHCP - Options.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-208	DHCP-Header	Description of the DHCPv4 packet structure	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		reserved; used to instantiate attributes dynamically as subscribers are referred		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		c e s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-209	DHCPv6-Header	Description of the DHCPv6 header	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		reserved; used to instantiate attributes dynamically as subscribers are referred		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		c e s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		stage.		that is, receive a packet of an ACP request from the

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			hexadecimal string	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			ic / T i m e d i n t e r v a l i n t e r i m 0 x 0 0 0 4 = I p a	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			ti	
			v	
			e	
			0	
			x	
			0	
			0	
			2	
			0	
			=	
			I	
			P	
			v	
			6	
			i	
			n	
			a	
			c	
			ti	
			v	
			e	
			0	
			x	
			0	
			0	
			4	
			0	
			=	
			S	
			e	
			s	
			s	
			i	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			0x000800=Session termination inactive	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			0x00000000	0 = Limited support 1 = Full support

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				significantly more than the change in age of the population

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				Processes signing of CoA request

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-211	Inner-Tag-Protocol-Id	Protocol identifier for the inner tag. The value is a hexadecimal number ranging from 0x0000 to 0xFFFF.	hexadecimal value ranging from 0x0000 to 0xFFFF	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			ff ff . • 0 x 8 1 0 0 = I n n e r V L A N t a g f o r d e s i g n a t	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				enabled
				2
				BSA
				sub
				scri
				bers

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-212	Routing-Services	Determines the minimum number of sessions with which the router can establish outgoing service connections	Integer: 4-octet • 0x00000000 = Disabled	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		bitly is enabled or disabled.	no for outgoing sessions.	
			• 00001 = Enable in	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			static allocation of resources using session server vices.	
			Any value other	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			0 or 1 is rejected.	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-213	Interface-Set-Targeting-Weight	Specify a weight for the interface	Integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		te it a n d it s m e m b e r l i n k s w i t h a n a g g r e g a t e d E t h		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Network member link for target aggregate distribution.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-214	Interface-Targeting-Weight	Specify a weight for the interface	integer : 4-octet	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		with an interface reference set and thus with the set's aggregator		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		egated Ethernet member link for target edge dist		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tribution. When an interface receives notification that it has a new weight		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		highest available for the first authentication		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		thorized subscribers in the network are used for these		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		e t.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-215	Queue-7-Output-Packets	(Junos OS Evolved) Number of packets sent over	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Queue 7 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		frame dissector.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-216	Hybrid-Access-DSL-Downstream-Speed	Specifies the downstream rate for the DSL leg	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		of a hybrid bridged access session tunnel endpoint for a subscriber. Used		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		yyt he PFE for load-balan cing traffic across the DS La		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		nd L T E le g s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-217	Hybrid-Access-LTE-Downstream-Speed	Specifies the downstream stream bandwidth and width for the LTE leg	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		of the hybrid access tunnel for subscribers.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dynamic by the Packet Forwarding Engine for load-balancing		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Enabling traffic across the SSLL and LTE g.s.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-218	Connection-Status-Message	Specifies connection parameters and settings	string	Yes

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		that is presented to other here most people/r/client (success)		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		assess home megabyte) . This is a logical extension to the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Reply Message attribute (18) and has the same		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		formats and semantics. The attributes described process users		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		only the first instance can receive multiple IP addresses		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		for this attribute.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-219	PON-Access-Type	Typology of PON transmission system in use: <ul style="list-style-type: none">0 - OTHER	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• 1 - GPON		
		• 2 - XG-PON1		
		• 3 - TWDM-PON		
		• 4 - XG		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			S - P O N	
		• 5 - W D M -		
			P O N	
		• 7 - U N K N O W N		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-220	ONT/ONU-Average-Data-Rate-Downstream	(PON) Average downstream rate for ONT	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		/ONU, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-221	ONT/ONU-Peak-Data-Rate-Downstream	(PON) Peak downstream rate for ONT/ONU	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		U , in K b p s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-222	ONT/ONU-Maximum-Data-Rate-Upstream	(PON) Maximum upstream data rate for ONT/ONU	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		U , in K b p s		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-223	ONT/ONU-Assured-Data-Rate-Upstream	(PON) Assured upstream data rate for ONT/O	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		NU, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-224	PON-Tree-Maximum-Data-Rate-Upstream	(PON) Maximum upstream data rate for PON	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tree, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-225	PON-Tree-Maximum-Data-Rate-Downstream	(PON) Maximum downstream rate	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ONtree, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-226	Expected-Throughput-Upstream	(Gigabit Ethernet expected upstream rate amount throughput, which is	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		the network data rate after the reduction by expected rate loss		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		s, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-227	Expected-Throughput-Downstream	(G.fast) Expected downstream rate amount throughput, which is	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		the network data rate after the reduction by expected rate loss		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		s, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-228	Attainable-Expected-Throughput-Upstream	(G.fast) Maximum attainable expected upstream rate	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		roughly put, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-229	Attainable-Expected-Throughput-Downstream	(Gigabit/sec) Maximum attainable expected downstream stream rate	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		throughput, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-230	Gamma-Data-Rate-Upstream	(G.fast) Actual upstream rate (net data rate)	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		e) for the local loop, adjust the down by any other group		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ut capacity limitation, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-231	Gamma-Data-Rate-Downstream	(G.fast) Actual downstream rate (net data rate)	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		attribute) for the local loop, adjust the down by any of the		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Input capacity limitation, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-232	Attainable-Gamma-Data-Rate-Upstream	(G.fast) Maximum attainable upstream data rate (n	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		extended attribute) for the local loop, adjust the down by an		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		throughput capability limitations, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-233	Attainable-Gamma-Data-Rate-Downstream	(G.f a st) M a x i m u m a t t a i n a b l e d o w n s t r e a m d a t a r a t e	3 2- b i t i n t e g e r	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable (network data attribute) for the local loop, adjust speed down by		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		any throughput capability limitation, in Kbps		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-239	Queue-0-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 0 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-240	Queue-0-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable OINTHECOURSE of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-241	Queue-1-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 1 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-242	Queue-1-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable intelligent use of distributed services to address		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		m e d u s e r.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-243	Queue-2-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 2 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-244	Queue-2-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 2 in the course of delivering this service to a fra		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-245	Queue-3-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 3 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-246	Queue-3-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 3 in the course of delivering this service to a fra		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-247	Queue-4-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 4 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-248	Queue-4-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 4 in the course of delivering this service to a fra		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-249	Queue-5-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 5 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-250	Queue-5-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 5 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-251	Queue-6-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 6 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-252	Queue-6-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 6 in the course of delivering this service to a fra		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		meduser.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-253	Queue-7-Output-Gigawords	(Junos OS Evolved) Number of times the Queue 7 o	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Output octets attributed to the router's 4-octet field.		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-254	Queue-7-Output-Octets	(Junos OS Evolved) Number of octets sent over Queue	32-bit integer	No

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enable 7 in the course of delivering this service to a		

Table 2: Supported Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	D e s c r i p t i o n	V a l u e	Dynamic CoA Support
		m e d u s e r.		

AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs

Table 3 on page 599 shows the RADIUS attributes and Juniper Networks VSAs (vendor ID 4874) support in AAA access messages. A checkmark in a column indicates that the message type supports that attribute.

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs

Attrib ute Numb er	Attribute Name	Access Request	Access Accept	Access Reject	Access Challen ge	CoA Reques t	Discon nect Reques t
1	User-Name	✓	✓	-	-	-	✓
2	User-Password	✓	-	-	-	-	-
3	CHAP-Password	✓	-	-	-	-	-
4	NAS-IP-Address	✓	-	-	-	-	-
5	NAS-Port	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
6	Service-Type	✓	✓	-	-	-	-
7	Framed-Protocol	✓	✓	-	-	-	-
8	Framed-IP-Address	✓	✓	-	-	✓	-
9	Framed-IP-Netmask	-	✓	-	-	-	-
11	Filter-Id	-	✓	-	-	-	-
12	Framed-MTU	✓	-	-	-	-	-
18	Reply-Message	-	✓	✓	✓	-	-
22	Framed-Route	-	✓	-	-	-	-
24	State	✓	✓	-	✓	-	-
25	Class	-	✓	-	-	✓	-
26-1	Virtual-Router	✓	✓	-	-	✓	-
26-4	Primary-DNS	-	✓	-	-	-	-
26-5	Secondary-DNS	-	✓	-	-	-	-
26-6	Primary-WINS	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-7	Secondary-WINS	-	✓	-	-	-	-
26-8	Tunnel-Virtual-Router	-	✓	-	-	-	-
26-9	Tunnel-Password	-	✓	-	-	-	-
26-10	Ingress-Policy-Name	-	✓	-	-	-	-
26-11	Egress-Policy-Name	-	✓	-	-	-	-
26-23	IGMP-Enable	-	✓	-	-	-	-
26-24	PPPoE-Description	✓	-	-	-	-	-
26-25	Redirect-VR-Name	-	✓	-	-	-	-
26-31	Service-Bundle	-	✓	-	-	-	-
26-33	Tunnel-Maximum-Sessions	-	✓	-	-	-	-
26-34	Framed-IP-Route-Tag. Supported only on JunosE for ERX and E320 platforms.	-	✓	-	-	-	-
26-47	Ipv6-Primary-DNS	-	✓	-	-	-	-
26-48	Ipv6-Secondary-DNS	-	✓	-	-	-	-
26-55	DHCP-Options	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-56	DHCP-MAC-Address	✓	✓	-	-	-	-
26-57	DHCP-GI-Address	✓	-	-	-	-	-
26-58	LI-Action	-	✓	-	-	✓	-
26-59	Med-Dev-Handle	-	✓	-	-	✓	-
26-60	Med-Ip-Address	-	✓	-	-	✓	-
26-61	Med-Port-Number	-	✓	-	-	✓	-
26-63	Interface-Desc	✓	-	-	-	-	-
26-64	Tunnel-Group	-	✓	-	-	-	-
26-65	Activate-Service	-	✓	-	-	✓	-
26-66	Deactivate-Service	-	✓	-	-	✓	-
26-67	Service-Volume	-	✓	-	-	✓	-
26-68	Service-Timeout	-	✓	-	-	✓	-
26-69	Service-Statistics	-	✓	-	-	✓	-
26-71	IGMP-Access-Name	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-72	IGMP-Access-Src-Name	-	✓	-	-	-	-
26-74	MLD-Access-Name	-	✓	-	-	-	-
26-75	MLD-Access-Src-Name	-	✓	-	-	-	-
26-77	MLD-Version	-	✓	-	-	-	-
26-78	IGMP-Version	-	✓	-	-	-	-
26-91	Tunnel-Switch-Profile	-	✓	-	-	-	-
26-92	L2C-Up-Stream-Data	✓	-	-	-	-	-
26-93	L2C-Down-Stream-Data	✓	-	-	-	-	-
26-94	Tunnel-Tx-Speed-Method	-	✓	-	-	-	-
26-97	IGMP-Immediate-Leave	-	✓	-	-	-	-
26-100	MLD-Immediate-Leave	-	✓	-	-	-	-
26-106	IPv6-Ingress-Policy-Name	-	✓	-	-	-	-
26-107	IPv6-Egress-Policy-Name	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-108	CoS-Parameter-Type	-	✓	-	-	✓	-
26-109	DHCP-Guided-Relay-Server	-	✓	-	-	-	-
26-110	Acc-Loop-Cir-Id	✓	-	-	-	-	-
26-111	Acc-Aggr-Cir-Id-Bin	✓	-	-	-	-	-
26-112	Acc-Aggr-Cir-Id-Asc	✓	-	-	-	-	-
26-113	Act-Data-Rate-Up	✓	-	-	-	-	-
26-114	Act-Data-Rate-Dn	✓	-	-	-	-	-
26-115	Min-Data-Rate-Up	✓	-	-	-	-	-
26-116	Min-Data-Rate-Dn	✓	-	-	-	-	-
26-117	Att-Data-Rate-Up	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-118	Att-Data-Rate-Dn	✓	-	-	-	-	-
26-119	Max-Data-Rate-Up	✓	-	-	-	-	-
26-120	Max-Data-Rate-Dn	✓	-	-	-	-	-
26-121	Min-LP-Data-Rate-Up	✓	-	-	-	-	-
26-122	Min-LP-Data-Rate-Dn	✓	-	-	-	-	-
26-123	Max-Interlv-Delay-Up	✓	-	-	-	-	-
26-124	Act-Interlv-Delay-Up	✓	-	-	-	-	-
26-125	Max-Interlv-Delay-Dn	✓	-	-	-	-	-
26-126	Act-Interlv-Delay-Dn	✓	-	-	-	-	-
26-127	DSL-Line-State	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-128	DSL-Type	✓	-	-	-	-	-
26-130	QoS-Set-Name	-	✓	-	-	-	-
26-140	Service-Interim-Account-Interval	-	✓	-	-	✓	-
26-141	Downstream-Calculated-QoS-Rate	✓	-	-	-	-	-
26-142	Upstream-Calculated-QoS-Rate	✓	-	-	-	-	-
26-143	Max-Clients-Per-Interface	-	✓	-	-	-	-
26-146	Cos-Scheduler-Pmt-Type	-	✓	-	-	✓	-
26-158	PPPoE-Padn	-	✓	-	-	-	-
26-159	Unisphere-L2tp-Csun-enable	-	✓	-	-	-	-
26-160	Vlan-Map-Id	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-161	IPv6-Delegated-Pool-Name	-	✓	-	-	-	-
26-162	Tx-Connect-Speed	✓	-	-	-	-	-
26-163	Rx-Connect-Speed	✓	-	-	-	-	-
26-164	IPv4-Release-Control	✓	-	-	-	-	-
26-173	Service-Activate-Type	-	✓	-	-	✓	-
26-174	Client-Profile-Name.	✓	✓	-	-	-	-
26-179	Service-Volume-Gigawords	-	✓	-	-	✓	-
26-180	Update-Service	-	-	-	-	✓	-
26-181	DHCPv6-Guided-Relay-Server	-	✓	-	-	-	-
26-182	Acc-Loop-Remote-Id	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-183	Acc-Loop-Encap	✓	-	-	-	-	-
26-184	Inner-Vlan-Map-Id	-	✓	-	-	-	-
26-189	DHCP-First-Relay-IPv4-Address	✓	-	-	-	-	-
26-190	DHCP-First-Relay-IPv6-Address	✓	-	-	-	-	-
26-191	Input-Interface-Filter	✓	-	-	-	✓	-
26-192	Output-Interface-Filter	✓	-	-	-	✓	-
26-193	Pim-Enable	-	✓	-	-	-	-
26-194	Bulk-CoA-Transaction-Id	-	-	-	-	✓	-
26-195	Bulk-CoA-Identifier	-	-	-	-	✓	-
26-196	IPv4-Input-Service-Set	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-197	IPv4-Output-Service-Set	✓	-	-	-	-	-
26-198	IPv4-Input-Service-Filter	✓	-	-	-	-	-
26-199	IPv4-Output-Service-Filter	✓	-	-	-	-	-
26-200	IPv6-Input-Service-Set	✓	-	-	-	-	-
26-201	IPv6-Output-Service-Set	✓	-	-	-	-	-
26-202	IPv6-Input-Service-Filter	✓	-	-	-	-	-
26-203	IPv6-Output-Service-Filter	✓	-	-	-	-	-
26-204	Adv-Pcef-Profile-Name	✓	-	-	-	-	-
26-205	Adv-Pcef-Rule-Name	✓	-	-	-	-	-
26-206	Re-Authentication-On-Renew	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-207	DHCPv6-Options	✓	✓	-	-	-	-
26-208	DHCP-Header	✓	-	-	-	-	-
26-209	DHCPv6-Header	✓	-	-	-	-	-
26-211	Inner-Tag-Protocol-Id	-	✓	-	-	-	-
26-212	Routing-Services	-	✓	-	-	-	-
26-213	Interface-Set-Targeting-Weight	-	✓	-	-	-	-
26-214	Interface-Targeting-Weight	-	✓	-	-	-	-
26-216	Hybrid-Access-DSL-Downstream-Speed	-	✓	-	-	-	-
26-217	Hybrid-Access-LTE-Downstream-Speed	-	✓	-	-	-	-
26-218	Connection-Status-Message	-	✓	-	-	✓	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-219	PON-Access-Type	✓	-	-	-	-	-
26-220	ONT/ONU-Average-Data-Rate-Downstream	✓	-	-	-	-	-
26-221	ONT/ONU-Peak-Data-Rate-Downstream	✓	-	-	-	-	-
26-222	ONT/ONU-Maximum-Data-Rate-Upstream	✓	-	-	-	-	-
26-223	ONT/ONU-Assured-Data-Rate-Upstream	✓	-	-	-	-	-
26-224	PON-Tree-Maximum-Data-Rate-Upstream	✓	-	-	-	-	-
26-225	PON-Tree-Maximum-Data-Rate-Downstream	✓	-	-	-	-	-
26-226	Expected-Throughput-Upstream	✓	-	-	-	-	-
26-227	Expected-Throughput-Downstream	✓	-	-	-	-	-
26-228	Attainable-Expected-Throughput-Upstream	✓	-	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-229	Attainable-Expected-Throughput-Downstream	✓	-	-	-	-	-
26-230	Gamma-Data-Rate-Upstream	✓	-	-	-	-	-
26-231	Gamma-Data-Rate-Downstream	✓	-	-	-	-	-
26-232	Attainable-Gamma-Data-Rate-Upstream	✓	-	-	-	-	-
26-233	Attainable-Gamma-Data-Rate-Downstream	✓	-	-	-	-	-
27	Session-Timeout	-	✓	-	✓	✓	-
28	Idle-Timeout	-	✓	-	✓	-	-
31	Calling-Station-ID	✓	-	-	-	✓	-
32	NAS-Identifier	✓	-	-	-	-	-
44	Acct-Session-ID	✓	-	-	-	✓	✓
61	NAS-Port-Type	✓	-	-	-	-	-
64	Tunnel-Type	✓	✓	-	-	-	-
65	Tunnel-Medium-Type	✓	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
66	Tunnel-Client-Endpoint	✓	✓	-	-	-	-
67	Tunnel-Server-Endpoint	✓	✓	-	-	-	-
68	Acct-Tunnel-Connection	✓	✓	-	-	-	-
69	Tunnel-Password	-	✓	-	-	-	-
82	Tunnel-Assignment-Id	✓	✓	-	-	-	-
83	Tunnel-Preference	-	✓	-	-	-	-
85	Acct-Interim-Interval	-	✓	-	-	-	-
87	NAS-Port-Id	✓	-	-	-	✓	-
88	Framed-Pool	-	✓	-	-	-	-
90	Tunnel-Client-Auth-Id	✓	✓	-	-	-	-
91	Tunnel-Server-Auth-Id	✓	✓	-	-	-	-
95	NAS-IPv6-Address	✓	-	-	-	-	-
96	Framed-Interface-ID	-	✓	-	-	-	-
97	Framed-IPv6-Prefix	-	✓	-	-	-	-

Table 3: AAA Access Messages: Supported RADIUS Attributes and Juniper Networks VSAs (Continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
98	Login-IPv6-Host	✓	✓	-	-	-	-
99	Framed-IPv6-Route	-	✓	-	-	-	-
100	Framed-IPv6-Pool	-	✓	-	-	-	-
123	Delegated-IPv6-Prefix	-	✓	-	-	-	-
168	Framed-IP-Address	-	✓	-	-	-	-
242	Ascend-Data-Filter	-	✓	-	-	✓	-



NOTE: The RADIUS VSA ID 26-174 denotes Client-Profile-Name. You can use RADIUS VSA 26-174 to override an assigned client dynamic profile with a specified client profile name. This VSA allows RADIUS to differentiate between dynamic profiles on the router when a version alias string is included. It supports both Access-Accept and Access-Request. This feature enhances profile management and customization capabilities within your network infrastructure.

AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs

Table 4 on page 615 shows the RADIUS attributes and Juniper Networks VSAs support in AAA accounting messages. A checkmark in a column indicates that the message type supports that attribute.

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
1	User-Name	✓	✓	✓	–	–
3	CHAP-Password	✓	–	–	–	–
4	NAS-IP-Address	✓	✓	✓	✓	✓
5	NAS-Port	✓	✓	✓	–	–
6	Service-Type	✓	✓	✓	–	–
7	Framed-Protocol	✓	✓	✓	–	–
8	Framed-IP-Address	✓	✓	✓	–	–
9	Framed-IP-Netmask	✓	✓	✓	–	–
11	Filter-Id	–	✓	✓	–	–
22	Framed-Route	✓	✓	✓	–	–
25	Class	✓	✓	✓	–	–
26-1	Virtual-Router	✓	✓	✓	–	–
26-10	Ingress-Policy-Name	✓	✓	✓	–	–
26-11	Egress-Policy-Name	✓	✓	✓	–	–
26-12	Policer-L1-Input-Gigawords	–	✓	✓	–	–

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-13	Policer-L1-Input-Octets	–	✓	✓	–	–
26-14	Policer-L2-Input-Gigawords	–	✓	✓	–	–
26-15	Policer-L2-Input-Octets	–	✓	✓	–	–
26-16	Policer-L3-Input-Gigawords	–	✓	✓	–	–
26-17	Policer-L3-Input-Octets	–	✓	✓	–	–
26-18	Policer-L4-Input-Gigawords	–	✓	✓	–	–
26-19	Policer-L4-Input-Octets	–	✓	✓	–	–
26-20	Policer-L1-Input-Gigapackets	–	✓	✓	–	–
26-21	Policer-L1-Input-Packets	–	✓	✓	–	–
26-24	PPPoE-Description	✓	✓	✓	–	–
26-29	Queue-0-Output-Gigapackets	–	✓	✓	–	–
26-32	Queue-0-Output-Packets	–	✓	✓	–	–
26-42	Input-Gigapackets	–	✓	✓	–	–
26-43	Output-Gigapackets	–	✓	✓	–	–
26-47	Ipv6-Primary-DNS	✓	✓	✓	–	–

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-48	Ipv6-Secondary-DNS	✓	✓	✓	-	-
26-49	Queue-1-Output-Gigapackets	-	✓	✓	-	-
26-50	Queue-1-Output-Packets	-	✓	✓	-	-
26-51	Disconnect-Cause	-	✓	-	-	-
26-55	DHCP-Options	✓	✓	✓	-	-
26-56	DHCP-MAC-Address	✓	✓	✓	-	-
26-57	DHCP-GI-Address	✓	✓	✓	-	-
26-63	Interface-Desc	✓	✓	✓	-	-
26-83	Service-Session	-	✓	✓	-	-
26-84	Policer-L2-Input-Gigapackets	-	✓	✓	-	-
26-85	Policer-L2-Input-Packets	-	✓	✓	-	-
26-86	Policer-L3-Input-Gigapackets	-	✓	✓	-	-
26-87	Policer-L3-Input-Packets	-	✓	✓	-	-
26-88	Policer-L3-Input-Gigapackets	-	✓	✓	-	-
26-89	Policer-L4-Input-Packets	-	✓	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-92	L2C-Up-Stream-Data	✓	✓	✓	–	–
26-93	L2C-Down-Stream-Data	✓	✓	✓	–	–
26-110	Acc-Loop-Cir-Id	✓	✓	✓	–	–
26-111	Acc-Aggr-Cir-Id-Bin	✓	✓	✓	–	–
26-112	Acc-Aggr-Cir-Id-Asc	✓	✓	✓	–	–
26-113	Act-Data-Rate-Up	✓	✓	✓	–	–
26-114	Act-Data-Rate-Dn	✓	✓	✓	–	–
26-115	Min-Data-Rate-Up	✓	✓	✓	–	–
26-116	Min-Data-Rate-Dn	✓	✓	✓	–	–
26-117	Att-Data-Rate-Up	✓	✓	✓	–	–
26-118	Att-Data-Rate-Dn	✓	✓	✓	–	–
26-119	Max-Data-Rate-Up	✓	✓	✓	–	–
26-120	Max-Data-Rate-Dn	✓	✓	✓	–	–
26-121	Min-LP-Data-Rate-Up	✓	✓	✓	–	–
26-122	Min-LP-Data-Rate-Dn	✓	✓	✓	–	–

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-123	Max-Interlv-Delay-Up	✓	✓	✓	-	-
26-124	Act-Interlv-Delay-Up	✓	✓	✓	-	-
26-125	Max-Interlv-Delay-Dn	✓	✓	✓	-	-
26-126	Act-Interlv-Delay-Dn	✓	✓	✓	-	-
26-127	DSL-Line-State	✓	✓	✓	-	-
26-128	DSL-Type	✓	✓	✓	-	-
26-131	Queue-2-Output-Gigapackets	-	✓	✓	-	-
26-132	Queue-2-Output-Packets	-	✓	✓	-	-
26-133	Queue-3-Output-Gigapackets	-	✓	✓	-	-
26-134	Queue-3-Output-Packets	-	✓	✓	-	-
26-135	Queue-4-Output-Gigapackets	-	✓	✓	-	-
26-136	Queue-4-Output-Packets	-	✓	✓	-	-
26-137	Queue-5-Output-Gigapackets	-	✓	✓	-	-
26-138	Queue-5-Output-Packets	-	✓	✓	-	-
26-139	Queue-6-Output-Gigapackets	-	✓	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-141	Downstream-Calculated-QoS-Rate	✓	✓	✓	-	-
26-142	Upstream-Calculated-QoS-Rate	✓	✓	✓	-	-
26-145	Queue-6-Output-Packets	-	✓	✓	-	-
26-149	Queue-7-Output-Gigapackets	-	✓	✓	-	-
26-151	IPv6-Acct-Input-Octets	-	✓	✓	-	-
26-152	IPv6-Acct-Output-Octets	-	✓	✓	-	-
26-153	IPv6-Acct-Input-Packets	-	✓	✓	-	-
26-154	IPv6-Acct-Output-Packets	-	✓	✓	-	-
26-155	IPv6-Acct-Input-Gigawords	-	✓	✓	-	-
26-156	IPv6-Acct-Output-Gigawords	-	✓	✓	-	-
26-160	Vlan-Map-Id	✓	✓	✓	-	-
26-162	Tx-Connect-Speed	✓	✓	✓	-	-
26-163	Rx-Connect-Speed	✓	✓	✓	-	-
26-164	IPv4-Release-Control	-	-	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-177	Cos-Shaping-Rate	✓	✓	✓	–	–
26-182	Acc-Loop-Remote-Id	✓	✓	–	–	–
26-183	Acc-Loop-Encap	✓	✓	–	–	–
26-184	Inner-Vlan-Map-Id	✓	✓	–	–	–
26-185	Core-Facing-Interface	✓	✓	–	–	–
26-188	DHCP-First-Relay-IPv4-Address	✓	✓	✓	–	–
26-190	DHCP-First-Relay-IPv6-Address	✓	✓	✓	–	–
26-191	Input-Interface-Filter	✓	✓	✓	–	–
26-192	Output-Interface-Filter	✓	✓	✓	–	–
26-207	DHCPv6-Options	✓	✓	✓	–	–
26-210	Acct-Request-Reason	✓	–	✓	–	–
26-215	Queue-7-Output-Packets	–	✓	✓	–	–
26-219	PON-Access-Type	✓	✓	✓	–	–
26-220	ONT/ONU-Average-Data-Rate-Downstream	✓	✓	✓	–	–

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-221	ONT/ONU-Peak-Data-Rate-Downstream	✓	✓	✓	-	-
26-222	ONT/ONU-Maximum-Data-Rate-Upstream	✓	✓	✓	-	-
26-223	ONT/ONU-Assured-Data-Rate-Upstream	✓	✓	✓	-	-
26-224	PON-Tree-Maximum-Data-Rate-Upstream	✓	✓	✓	-	-
26-225	PON-Tree-Maximum-Data-Rate-Downstream	✓	✓	✓	-	-
26-226	Expected-Throughput-Upstream	✓	✓	✓	-	-
26-227	Expected-Throughput-Downstream	✓	✓	✓	-	-
26-228	Attainable-Expected-Throughput-Upstream	✓	✓	✓	-	-
26-229	Attainable-Expected-Throughput-Downstream	✓	✓	✓	-	-
26-230	Gamma-Data-Rate-Upstream	✓	✓	✓	-	-
26-231	Gamma-Data-Rate-Downstream	✓	✓	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-232	Attainable-Gamma-Data-Rate-Upstream	✓	✓	✓	-	-
26-233	Attainable-Gamma-Data-Rate-Downstream	✓	✓	✓	-	-
26-239	Queue-0-Output-Gigawords	-	✓	✓	-	-
26-240	Queue-0-Output-Octets	-	✓	✓	-	-
26-241	Queue-1-Output-Gigawords	-	✓	✓	-	-
26-242	Queue-1-Output-Octets	-	✓	✓	-	-
26-243	Queue-2-Output-Gigawords	-	✓	✓	-	-
26-244	Queue-2-Output-Octets	-	✓	✓	-	-
26-245	Queue-3-Output-Gigawords	-	✓	✓	-	-
26-246	Queue-3-Output-Octets	-	✓	✓	-	-
26-247	Queue-4-Output-Gigawords	-	✓	✓	-	-
26-248	Queue-4-Output-Octets	-	✓	✓	-	-
26-249	Queue-5-Output-Gigawords	-	✓	✓	-	-
26-250	Queue-5-Output-Octets	-	✓	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-251	Queue-6-Output-Gigawords	-	✓	✓	-	-
26-252	Queue-6-Output-Octets	-	✓	✓	-	-
26-253	Queue-7-Output-Gigawords	-	✓	✓	-	-
26-254	Queue-7-Output-Octets	-	✓	✓	-	-
31	Calling-Station-ID	✓	✓	✓	-	-
32	NAS-Identifier	✓	✓	✓	✓	✓
40	Acct-Status-Type	✓	✓	✓	✓	✓
41	Acct-Delay-Time	✓	✓	✓	✓	✓
42	Acct-Input-Octets	-	✓	✓	-	-
43	Acct-Output-Octets	-	✓	✓	-	-
44	Acct-Session-ID	✓	✓	✓	✓	✓
45	Acct-Authentic	✓	✓	✓	✓	✓
46	Acct-Session-Time	-	✓	✓	-	-
47	Acct-Input-Packets	-	✓	✓	-	-
48	Acct-Output-Packets	-	✓	✓	-	-

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
49	Acct-Terminate-Cause	–	✓	✓	–	–
52	Acct-Input-Gigawords	–	✓	✓	–	–
53	Acct-Output-Gigawords	–	✓	✓	–	–
55	Event-Timestamp	✓	✓	✓	✓	✓
61	NAS-Port-Type	✓	✓	✓	–	–
64	Tunnel-Type	✓	✓	✓	–	–
65	Tunnel-Medium-Type	✓	✓	✓	–	–
66	Tunnel-Client-Endpoint	✓	✓	✓	–	–
67	Tunnel-Server-Endpoint	✓	✓	✓	–	–
68	Acct-Tunnel-Connection	✓	✓	✓	–	–
77	Connect-Info	✓	✓	–	–	–
82	Tunnel-Assignment-Id	✓	✓	✓	–	–
87	NAS-Port-Id	✓	✓	✓	–	–
90	Tunnel-Client-Auth-Id	✓	✓	✓	–	–
91	Tunnel-Server-Auth-Id	✓	✓	✓	–	–

Table 4: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs
(Continued)

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
99	Framed-IPv6-Route	✓	✓	✓	–	–
100	Framed-IPv6-Pool	✓	✓	✓	–	–
123	Delegated-IPv6-Prefix	✓	✓	✓	–	–

DSL Forum Vendor-Specific Attributes

IN THIS SECTION

- [DSL Forum VSAs and PPPoE-IA Tags | 634](#)

Broadband access lines have many characteristics that are not supported by standard RADIUS attributes. A telecommunications and networking industry consortium, formerly called the DSL Forum and since 2008 called the Broadband Forum, develops standards and specifications for broadband technologies and products. The DSL Forum concentrated only on digital subscriber lines. The forum changed its name as it expanded the scope of its work to other broadband access technologies, such as passive optical networking (PON).

The DSL Forum defined RADIUS vendor-specific attributes (VSAs) to convey that information to the RADIUS server for processing. These VSAs include information about the access lines, the subscribers using the lines, and data rates on the lines. Subscriber management does not process the VSA values—the router simply passes the values received from the subscriber to the RADIUS server, without performing any parsing or manipulation. However, you can manage the content of the VSAs either by using the client configuration to restrict the DSL Forum VSAs that the client sends, or by configuring the RADIUS server to ignore unwanted DSL Forum VSAs.

The terminology used with the DSL Forum VSAs can be confusing. Each of these VSAs is actually a subattribute of the DSL Forum RADIUS VSA. The DSL Forum RADIUS VSA is simply a container for the subattributes that transports them to the RADIUS server. The DSL Forum RADIUS VSA provides the following information that applies to each subattribute:

- Type = 26. This value indicates that the subattribute is a vendor-specific attribute.
- Vendor-ID = 3561. This value is the vendor ID (enterprise number) assigned to the Broadband Forum by the Internet Assigned Numbers Authority (IANA).

Each subattribute is a TLV; that is, it specifies type, length, and value information:

- The vendor type is a number assigned by the Broadband Forum that identifies the subattribute. This number is sometimes referred to as the attribute number.
- The vendor length is a number that specifies the length of the entire subattribute.
- The value field contains information specific to the subattribute, such as data rates or access line identifiers.

After the name changed to the Broadband Forum, the forum added PON VSAs. We still refer to them as DSL Forum VSAs because they are subattributes of the DSL Forum VSA. Some of the VSAs previously used only for DSL networks are also used for PON networks.



NOTE: The full designation for a DSL Forum VSA is 26-3561-*type*. The vendor ID is critical to distinguishing between VSAs. For example, 26-3561-1 is a different attribute than 26-4874-1; 4874 is a Juniper Networks enterprise number. When the enterprise is clear from the context, our documentation may omit the enterprise number. For example, when a table refers to attributes for only one enterprise, we may omit the number to make the table easier to read.

The following documents provide information about the attributes:

- 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*

[Table 5 on page 628](#) describes the DSL Forum VSAs. Starting in Junos OS Release 19.3R1, we support the PON and DSL G.fast VSAs.

Table 5: DSL Forum VSAs (Vendor ID 3561)

Type	Name	Description	Access Type	Value
1	Agent-Circuit-Id	<p>Identifier for the subscriber agent circuit ID (ACI) that corresponds to the access node interface from which subscriber requests are initiated.</p> <p>For auto-sensed VLANs, the ACI is extracted from DHCP discover, DHCPv6 solicit, or PPPoE PADI messages, stored in the VLAN shared database entry, and then presented in the RADIUS Access-Request message in this VSA.</p>	DSL, PON	string
2	Agent-Remote-Id	<p>Unique identifier for the subscriber associated with the access node interface from which requests are initiated.</p> <p>For auto-sensed VLANs, the ARI is extracted from DHCP discover, DHCPv6 solicit, or PPPoE PADI messages, stored in the VLAN shared database entry, and then presented in the RADIUS Access-Request message in this VSA.</p>	DSL, PON	string
3	Access-Aggregation-Circuit-ID-ASCII	<p>ASCII identifier for the subscriber access line, based on its network-facing logical appearance</p> <p>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.</p>	DSL, PON	string

Table 5: DSL Forum VSAs (Vendor ID 3561) (Continued)

Type	Name	Description	Access Type	Value
6	Access-Aggregation-Circuit-ID-Binary	Binary identifier for the subscriber access line	DSL, PON	string
129	Actual-Data-Rate-Upstream	Actual upstream data rate of the subscriber's synchronized DSL link, in bps	DSL	32-bit integer
130	Actual-Data-Rate-Downstream	Actual downstream data rate of the subscriber's synchronized DSL link, in bps	DSL	32-bit integer
131	Minimum-Data-Rate-Upstream	Minimum upstream data rate configured for the subscriber, in bps	DSL	32-bit integer
132	Minimum-Data-Rate-Downstream	Minimum downstream data rate configured for the subscriber, in bps	DSL	32-bit integer
133	Attainable-Data-Rate-Upstream	Upstream data rate that the subscriber can attain, in bps	DSL	32-bit integer
134	Attainable-Data-Rate-Downstream	Downstream data rate that the subscriber can attain, in bps	DSL	32-bit integer
135	Maximum-Data-Rate-Upstream	Maximum upstream data rate configured for the subscriber, in bps	DSL	32-bit integer
136	Maximum-Data-Rate-Downstream	Maximum downstream data rate configured for the subscriber, in bps	DSL	32-bit integer
137	Minimum-Data-Rate-Upstream-Low-Power	Minimum upstream data rate in low power state configured for the subscriber, in bps	DSL	32-bit integer

Table 5: DSL Forum VSAs (Vendor ID 3561) (Continued)

Type	Name	Description	Access Type	Value
138	Minimum-Data-Rate-Downstream-Low-Power	Minimum downstream data rate in low power state configured for the subscriber, in bps	DSL	32-bit integer
139	Maximum-Interleaving-Delay-Upstream	Maximum one-way upstream interleaving delay configured for the subscriber, in milliseconds	DSL	32-bit integer
140	Actual-Interleaving-Delay-Upstream	Subscriber's actual one-way upstream interleaving delay, in milliseconds	DSL	32-bit integer
141	Maximum-Interleaving-Delay-Downstream	Maximum one-way downstream interleaving delay configured for the subscriber, in milliseconds	DSL	32-bit integer
142	Actual-Interleaving-Delay-Downstream	Subscriber's actual one-way downstream interleaving delay, in milliseconds	DSL	32-bit integer
144	Access-Loop-Encapsulation	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated	DSL, PON	string: 3-byte

Table 5: DSL Forum VSAs (Vendor ID 3561) *(Continued)*

Type	Name	Description	Access Type	Value
145	DSL-Type	<p>Type of DSL transmission system in use:</p> <ul style="list-style-type: none"> • 0—OTHER • 1—ADSL1 • 2—ADSL2 • 3—ADSL2+ • 4—VDSL1 • 5—VDSL2 • 6—SDSL • 8—G.fast • 9—VDSL2 Annex Q • 10—SDSL bonded • 11—VDSL2 bonded • 12—G.fast bonded • 13—VDSL2 Annex Q bonded 	DSL	32-bit integer

Table 5: DSL Forum VSAs (Vendor ID 3561) (Continued)

Type	Name	Description	Access Type	Value
146	PON-Access-Type	Type of PON transmission system in use: <ul style="list-style-type: none"> • 0—OTHER • 1—GPON • 2—XG-PON1 • 3—TWDM-PON • 4—XGS-PON • 5—WDM-PON • 7—UNKNOWN 	PON	32-bit integer
147	ONT/ONU-Average-Data-Rate-Downstream	Average downstream data rate for ONT/ONU, in Kbps	PON	32-bit integer
148	ONT/ONU-Peak-Data-Rate-Downstream	Peak downstream data rate for ONT/ONU, in Kbps	PON	32-bit integer
149	ONT/ONU-Maximum-Data-Rate-Upstream	Maximum upstream data rate for ONT/ONU, in Kbps	PON	32-bit integer
150	ONT/ONU-Assured-Data-Rate-Upstream	Assured upstream data rate for ONT/ONU, in Kbps	PON	32-bit integer
151	PON-Tree-Maximum-Data-Rate-Upstream	Maximum upstream data rate for the PON tree, in Kbps	PON	32-bit integer
152	PON-Tree-Maximum-Data-Rate-Downstream	Maximum downstream data rate for the PON tree, in Kbps	PON	32-bit integer

Table 5: DSL Forum VSAs (Vendor ID 3561) (Continued)

Type	Name	Description	Access Type	Value
155	Expected-Throughput-Upstream	Expected upstream throughput, which is the net data rate reduced by expected rate loss, in Kbps	G.fast (DSL)	32-bit integer
156	Expected-Throughput-Downstream	Expected upstream throughput, which is the net data rate reduced by expected rate loss, in Kbps	G.fast (DSL)	32-bit integer
157	Attainable-Expected-Throughput-Upstream	Maximum attainable expected upstream throughput, in Kbps	G.fast (DSL)	32-bit integer
158	Attainable-Expected-Throughput-Downstream	Maximum attainable expected downstream throughput, in Kbps	G.fast (DSL)	32-bit integer
159	Gamma-Data-Rate-Upstream	Actual upstream data rate (net data rate) for the local loop, adjusted down by any throughput capability limitations, in Kbps	G.fast (DSL)	32-bit integer
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	G.fast (DSL)	32-bit integer
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 Kbps	G.fast (DSL)	32-bit integer
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 Kbps	G.fast (DSL)	32-bit integer

Table 5: DSL Forum VSAs (Vendor ID 3561) (Continued)

Type	Name	Description	Access Type	Value
254	IWF-Session	Indication that the interworking function (IWF) has been performed for the subscriber's PPPoA over PPPoE session	DSL	No data field required

DSL Forum VSAs and PPPoE-IA Tags

In addition to using information received in ANCP messages, the ANCP agent on the router can use access line information conveyed in PPPoE packets, such as the PADI and PADR discovery packets. For PPPoE subscribers that connect through an access node that is running ANCP, the access node adds access-line information to PPPoE intermediate agent (PPPoE-IA) tags. These tags are located in the discovery packets that it passes to the router during the establishment of dynamic PPPoE sessions. Similarly to the way access line information is carried in sub-attributes of the DSL Forum VSA, this information is contained in sub-tags in the PPPoE Vendor-Specific-Tag (0x105). The sub-tags are also called tags. The data represents a current, accurate snapshot of the values at the moment that the subscriber connection is initiated.

[Table 6 on page 634](#) shows the PPPoE-IA tags that correspond to the DSL Forum VSAs. The tag value is simply the hexadecimal equivalent of the VSA type number. The vendor ID is the same for both the DSL Forum VSAs and the PPPoE tags: 3561 (0xDE9).

Table 6: Correlation Between DSL Forum VSAs and PPPoE-IA Tags

VSA Type	VSA Name	PPPoE Tag
1	Agent-Circuit-Id	0x01
2	Agent-Remote-Id	0x02
3	Access-Aggregation-Circuit-ID-ASCII	0x03
6	Access-Aggregation-Circuit-ID-Binary	0x06
129	Actual-Data-Rate-Upstream	0x81

Table 6: Correlation Between DSL Forum VSAs and PPPoE-IA Tags *(Continued)*

VSA Type	VSA Name	PPPoE Tag
130	Actual-Data-Rate-Downstream	0x82
131	Minimum-Data-Rate-Upstream	0x83
132	Minimum-Data-Rate-Downstream	0x84
133	Attainable-Data-Rate-Upstream	0x85
134	Attainable-Data-Rate-Downstream	0x86
135	Maximum-Data-Rate-Upstream	0x87
136	Maximum-Data-Rate-Downstream	0x88
137	Minimum-Data-Rate-Upstream-Low-Power	0x89
138	Minimum-Data-Rate-Downstream-Low-Power	0x8A
139	Maximum-Interleaving-Delay-Upstream	0x8B
140	Actual-Interleaving-Delay-Upstream	0x8C
141	Maximum-Interleaving-Delay-Downstream	0x8D
142	Actual-Interleaving-Delay-Downstream	0x8D
144	Access-Loop-Encapsulation	0x90
145	DSL-Type	0x91

Table 6: Correlation Between DSL Forum VSAs and PPPoE-IA Tags (Continued)

VSA Type	VSA Name	PPPoE Tag
146	PON-Access-Type	0x92
147	ONT/ONU-Average-Data-Rate-Downstream	0x93
148	ONT/ONU-Peak-Data-Rate-Downstream	0x94
149	ONT/ONU-Maximum-Data-Rate-Upstream	0x95
150	ONT/ONU-Assured-Data-Rate-Upstream	0x96
151	PON-Tree-Maximum-Data-Rate-Upstream	0x97
152	PON-Tree-Maximum-Data-Rate-Downstream	0x98
155	Expected-Throughput-Upstream	0x9B
156	Expected-Throughput-Downstream	0x9C
157	Attainable-Expected-Throughput-Upstream	0x9D
158	Attainable-Expected-Throughput-Downstream	0x9E
159	Gamma-Data-Rate-Upstream	0x9F
160	Gamma-Data-Rate-Downstream	0xA0
161	Attainable-Gamma-Data-Rate-Upstream	0xA1
162	Attainable-Gamma-Data-Rate-Downstream	0xA2

Table 6: Correlation Between DSL Forum VSAs and PPPoE-IA Tags (Continued)

VSA Type	VSA Name	PPPoE Tag
254	IWF-Session	0xFE

DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS

Table 7 on page 637 lists the DSL Forum VSAs supported by Junos OS in RADIUS Access-Request, Acct-Start, Acct-Stop, Interim-Acct, and CoA-Request messages. A checkmark in a column indicates that the message type supports that attribute.



NOTE: The DSL Forum vendor ID is 3561 is omitted from the attribute number to simplify the table. For example, the full designation for DSL Forum VSA Agent-Circuit-Id is 26-3561-1.

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561)

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-1	Agent-Circuit-Id	✓	✓	✓	✓	✓
26-2	Agent-Remote-Id	✓	✓	✓	✓	✓
26-3	Access-Aggregation-Circuit-ID-ASCII	✓	✓	✓	✓	–
26-6	Access-Aggregation-Circuit-ID-Binary	✓	✓	✓	✓	–

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561) *(Continued)*

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-129	Actual-Data-Rate-Upstream	✓	✓	✓	✓	–
26-130	Actual-Data-Rate-Downstream	✓	✓	✓	✓	–
26-131	Minimum-Data-Rate-Upstream	✓	✓	✓	✓	–
26-132	Minimum-Data-Rate-Downstream	✓	✓	✓	✓	–
26-133	Attainable-Data-Rate-Upstream	✓	✓	✓	✓	–
26-134	Attainable-Data-Rate-Downstream	✓	✓	✓	✓	–
26-135	Maximum-Data-Rate-Upstream	✓	✓	✓	✓	–
26-136	Maximum-Data-Rate-Downstream	✓	✓	✓	✓	–

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561) (Continued)

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-137	Minimum-Data-Rate-Upstream-Low-Power	✓	✓	✓	✓	-
26-138	Minimum-Data-Rate-Downstream-Low-Power	✓	✓	✓	✓	-
26-139	Maximum-Interleaving-Delay-Upstream	✓	✓	✓	✓	-
26-140	Actual-Interleaving-Delay-Upstream	✓	✓	✓	✓	-
26-141	Maximum-Interleaving-Delay-Downstream	✓	✓	✓	✓	-
26-142	Actual-Interleaving-Delay-Downstream	✓	✓	✓	✓	-
26-144	Access-Loop-Encapsulation	✓	✓	✓	✓	-
26-145	DSL-Type	✓	✓	✓	✓	-

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561) *(Continued)*

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-146	PON-Access-Type	✓	✓	✓	✓	–
26-147	ONT/ONU-Average-Data-Rate-Downstream	✓	✓	✓	✓	–
26-148	ONT/ONU-Peak-Data-Rate-Downstream	✓	✓	✓	✓	–
26-149	ONT/ONU-Maximum-Data-Rate-Upstream	✓	✓	✓	✓	–
26-150	ONT/ONU-Assured-Data-Rate-Upstream	✓	✓	✓	✓	–
26-151	PON-Tree-Maximum-Data-Rate-Upstream	✓	✓	✓	✓	–
26-152	PON-Tree-Maximum-Data-Rate-Downstream	✓	✓	✓	✓	–

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561) *(Continued)*

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-155	Expected-Throughput-Upstream	✓	✓	✓	✓	–
26-156	Expected-Throughput-Downstream	✓	✓	✓	✓	–
26-157	Attainable-Expected-Throughput-Downstream	✓	✓	✓	✓	–
26-158	Attainable-Expected-Throughput-Downstream	✓	✓	✓	✓	–
26-159	Gamma-Data-Rate-Upstream	✓	✓	✓	✓	–
26-160	Gamma-Data-Rate-Downstream	✓	✓	✓	✓	–
26-161	Attainable-Gamma-Data-Rate-Upstream	✓	✓	✓	✓	–
26-162	Attainable-Gamma-Data-Rate-Downstream	✓	✓	✓	✓	–

Table 7: RADIUS Message Support for DSL Forum VSAs (Vendor ID 3561) (Continued)

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
26-254	IWF-Session	✓	✓	✓	✓	–

RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses

Starting in Junos OS Release 15.1, the Junos OS AAA implementation supports RADIUS VSAs that identify the primary and secondary DNS servers for IANA private enterprise number 311 (Microsoft Corporation). For example, during PPP authentication, the router receives the VSAs from a RADIUS server and uses the attributes to provision customer premise equipment.

The two VSAs are shown in the following table, and are described in RFC 2548 (*Microsoft Vendor-specific RADIUS Attributes*)

Table 8: Microsoft Vendor-Specific RADIUS Attributes for DNS Server Addresses

Attribute Number	Attribute Name	Description	Value
26-28	MS-Primary-DNS-Server	IP address of the primary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.	integer: 4-octet <i>primary-dns-address</i>
26-29	MS-Secondary-DNS-Server	IP address of the secondary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.	integer: 4-octet <i>secondary-dns-address</i>

SEE ALSO

[DNS Address Assignment Precedence](#) | **1106**

Support for Cisco Systems VSAs

Cisco Systems, IANA private enterprise number 9, uses a single VSA, Cisco-AVPair (26-1). This VSA conveys different information based on the values it contains. In some subscriber access networks, which have a Junos based BNG connected to both a RADIUS server and a Cisco BroadHop application that is used as the Policy Control and Charging Rules Function (PCRF) server for provisioning services using RADIUS change of authorization (CoA) messages, you can use this VSA in RADIUS messages to activate and deactivate services. You cannot modify any attributes in authentication, accounting, or CoA responses in the RADIUS messages that the BNG sends. See ["Processing Cisco VSAs in RADIUS Messages for Service Provisioning" on page 754](#) for more information.

Any Cisco VSAs other than the ones used to provision the services are considered as unsupported attributes.

Subscriber Management RADIUS Dictionary Files

The Juniper Networks RADIUS dictionary that is used by default for subscriber management is updated when software features that affect the file are added or changed. The dictionary is not updated for every Junos OS release. The dictionary includes Juniper Networks vendor-specific attributes that are used by Junos OS, JunosE OS, or both.



NOTE: The VSA names in the dictionary begin with the prefix "Jnpr-" or "Unisphere". By convention, both prefixes are omitted from the Tech Library documentation to reduce confusion in feature discussions.

- ["Junos OS Release 18.4 Subscriber Management RADIUS Dictionary \[DCT\]" on page 990](#)
- ["Junos OS Release 18.2 Subscriber Management RADIUS Dictionary \[DCT\]" on page 968](#)
- ["Junos OS Release 17.4 Subscriber Management RADIUS Dictionary \[DCT\]" on page 946](#)
- ["Junos OS Release 17.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 924](#)
- ["Junos OS Release 16.2 Subscriber Management RADIUS Dictionary \[DCT\]" on page 902](#)
- ["Junos OS Release 16.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 881](#)
- ["Junos OS Release 15.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 859](#)

Interface Text Descriptions for Inclusion in RADIUS Attributes

RADIUS attributes such as NAS-Port-ID (87) and Calling-Station-ID (31) include a description that identifies the physical interface that is used to authenticate subscribers. The default format for nonchannelized interfaces is as follows:

```
interface-type-slot/adapter/port.subinterface[:svlan-vlan]
```

For example, consider physical interface ge-1/2/0, with a subinterface of 100 and SVLAN identifier of 100. The interface description used in the NAS-Port-ID is ge-1/2/0.100:100.

Starting in Junos OS Release 17.3R1, a different format is used for channelized interfaces. For channelized interfaces, the default interface description is as follows:

```
interface-type-slot/adapter/logical-port-number.subinterface[:svlan-vlan]
```

The channel information (logical port number) is determined by this formula:

Logical port number = 100 + (actual-port-number x 20) + channel-number

For example, consider a channelized interface 3 on port 2 where the:

- Physical interface is xe-0/1/2:3.
- Subinterface is 4.
- SVLAN is 5.
- VLAN is 6.

Using the formula, the logical port number = 100 + (2 x 20) + 3 = 143. Consequently, the default interface description is xe-0/1/143.4-5.6.

You can optionally configure the interface description format in an access profile to exclude the adapter, channel, or subinterface information.

For example, if you exclude the subinterface from the nonchannelized interface description format, the description becomes ge-1/2/0:100. If you exclude the channel information from the channelized interface description format, the description becomes xe-0/1/2.4-5.6.

SEE ALSO

RADIUS Servers and Parameters for Subscriber Access 646
Configuring a Calling-Station-ID with Additional Options 660

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, a different format is used for channelized interfaces.

- | | |
|------|---|
| 15.1 | Starting in Junos OS Release 15.1, the Junos OS AAA implementation supports RADIUS VSAs that identify the primary and secondary DNS servers for IANA private enterprise number 311 (Microsoft Corporation). |
|------|---|

RELATED DOCUMENTATION

[RADIUS Authentication and Accounting Basic Configuration | 721](#)

[RADIUS NAS Port Attributes and Options | 689](#)

Subscriber Management RADIUS Dictionary Files

The Juniper Networks RADIUS dictionary that is used by default for subscriber management is updated when software features that affect the file are added or changed. The dictionary is not updated for every Junos OS release. The dictionary includes Juniper Networks vendor-specific attributes that are used by Junos OS, JunosE OS, or both.



NOTE: The VSA names in the dictionary begin with the prefix “Jnpr-” or “Unisphere”. By convention, both prefixes are omitted from the Tech Library documentation to reduce confusion in feature discussions.

- ["Junos OS Release 18.4 Subscriber Management RADIUS Dictionary \[DCT\]" on page 990](#)
- ["Junos OS Release 18.2 Subscriber Management RADIUS Dictionary \[DCT\]" on page 968](#)
- ["Junos OS Release 17.4 Subscriber Management RADIUS Dictionary \[DCT\]" on page 946](#)
- ["Junos OS Release 17.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 924](#)
- ["Junos OS Release 16.2 Subscriber Management RADIUS Dictionary \[DCT\]" on page 902](#)
- ["Junos OS Release 16.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 881](#)
- ["Junos OS Release 15.1 Subscriber Management RADIUS Dictionary \[DCT\]" on page 859](#)

RADIUS for Subscriber Management

IN THIS CHAPTER

- [RADIUS Servers and Parameters for Subscriber Access | 646](#)
- [Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers | 668](#)
- [Session Options for Subscriber Access | 674](#)
- [RADIUS NAS Port Attributes and Options | 689](#)
- [RADIUS Logical Line Identification | 714](#)
- [RADIUS Authentication and Accounting Basic Configuration | 721](#)
- [RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 727](#)
- [Configuring RADIUS Reauthentication for DHCP Subscribers | 738](#)
- [RADIUS Accounting for Subscriber Access | 741](#)
- [Verifying and Managing Subscriber AAA Information | 772](#)
- [Session Termination Causes and RADIUS Termination Cause Codes | 774](#)
- [AAA Termination Causes and Code Values | 779](#)
- [DHCP Termination Causes and Code Values | 781](#)
- [L2TP Termination Causes and Code Values | 783](#)
- [PPP Termination Causes and Code Values | 809](#)
- [VLAN Termination Causes and Code Values | 822](#)

RADIUS Servers and Parameters for Subscriber Access

IN THIS SECTION

- [RADIUS Authentication and Accounting Server Definition | 647](#)
- [Configuring Options that Apply to All RADIUS Servers | 650](#)

- [Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | 652](#)
- [Configuring Access Profile Options for Interactions with RADIUS Servers | 653](#)
- [Configuring a Calling-Station-ID with Additional Options | 660](#)
- [Filtering RADIUS Attributes and VSAs from RADIUS Messages | 665](#)

Configuring parameters and options for RADIUS servers is a major part of your subscriber management configuration. After defining the authentication and accounting servers, you configure options for all RADIUS servers. You also configure access profiles that enable you to specify subscriber access authentication, authorization and accounting configuration parameters for subscribers or groups of subscribers. The profile settings override global settings. Although some options are available at both the global level and the access profile level, many options are available only in access profiles.

After you have created an access profile, you must specify where the profile is used with an access-profile statement; this is known as attaching the profile. Access profiles can be assigned at various levels. For example, some of places you can attach access profiles

- Globally for a routing instance.
- In dynamic profiles.
- In a domain map, which maps access options and session parameters for subscriber sessions.
- On the interfaces for dynamic VLANs and dynamic stacked VLANs.
- On the interface or in a subscriber group for subscribers with statically configured interfaces for dynamic service provisioning.
- On DHCP relay agents and DHCP local servers for DHCP clients or subscribers.

Because you can attach access profiles at many levels, the most specific access profile takes precedence over any other profile assignments to avoid conflict. Authentication and accounting do not run unless you attach the profile.

RADIUS Authentication and Accounting Server Definition

When you use RADIUS for subscriber management, you must define one or more external RADIUS servers that the router communicates with for subscriber authentication and accounting. Besides specifying the IPv4 or IPv6 address of the server, you can configure options and attributes that determine how the router interacts with the specified servers.

You can define RADIUS servers and connectivity options at the `[edit access radius-server]` hierarchy level, at the `[edit access profile name radius-server]` hierarchy level, or at both levels.



NOTE: The AAA process (authd) determines which server definitions to use as follows:

- When RADIUS server definitions are present only in `[edit access radius-server]`, authd uses those definitions.
- When RADIUS server definitions are present only in the access profile, authd uses those definitions.
- When RADIUS server definitions are present in both `[edit access radius-server]` and in the access profile, authd uses only the access profile definitions.

To use a RADIUS server, you must designate it as an authentication server, an accounting server, or both, in an access profile. You must do so for servers regardless of whether they are defined in an access profile or at the `[edit access radius-server]` hierarchy level.

To define RADIUS servers and to specify how the router interacts with the server:



NOTE: This procedure shows only the `[edit access radius-server]` hierarchy level. You can optionally configure any of these parameters at the `[edit access profile profile-name radius-server]` hierarchy level. You can do so either in addition to the global setting or instead of the global setting. When you apply a profile, the profile settings override the global configuration.

1. Specify the IPv4 or IPv6 address of the RADIUS server.

```
[edit access]
user@host# edit radius-server server-address
```

2. (Optional) Configure the RADIUS server accounting port number.

```
[edit access radius-server server-address]
user@host# set accounting-port port-number
```

3. (Optional) Configure the port number the router uses to contact the RADIUS server.

```
[edit access radius-server server-address]
user@host# set port port-number
```

4. Configure the required secret (password) that the local router passes to the RADIUS client. Secrets enclosed in quotation marks can contain spaces.

```
[edit access radius-server server-address]  
user@host# set secret password
```

5. (Optional) Configure the maximum number of outstanding requests that a RADIUS server can maintain. An outstanding request is a request to which the RADIUS server has not yet responded.

```
[edit access radius-server server-address]  
user@host# set max-outstanding-requests value
```

6. Configure the source address for the RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 or IPv6 address configured on one of the router interfaces.

```
[edit access radius-server server-address]  
user@host# set source-address source-address
```

7. (Optional) Configure retry and timeout values for authentication and accounting messages.
 - a. Configure how many times the router attempts to contact a RADIUS server when it has received no response.

```
[edit access radius-server server-address]  
user@host# set retry number
```

- b. Configure how long the router waits to receive a response from a RADIUS server before retrying the contact.

```
[edit access radius-server server-address]  
user@host# set timeout seconds
```



NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.



NOTE: The retry and timeout settings apply to both authentication and accounting messages unless you configure both the `accounting-retry` statement and the `accounting-timeout` statement. In that case, the retry and timeout settings apply only to authentication messages.

8. (Optional) Configure retry and timeout values for accounting messages separate from the settings for authentication messages.



NOTE: You must configure both the `accounting-retry` and the `accounting-timeout` statements. If you do not, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

- a. Configure how many times the router attempts to send accounting messages to the RADIUS accounting server when it has received no response.

```
[edit access radius-server server-address]
user@host# set accounting-retry number
```

- b. Configure how long the router waits to receive a response from a RADIUS accounting server before retrying the request.

```
[edit access radius-server server-address]
user@host# set accounting-timeout seconds
```

9. (Optional) Configure the router to contact the RADIUS server for logical line identification (LLID) preauthentication requests. See ["RADIUS Logical Line Identification" on page 714](#).
10. (Optional) Configure the port that the router monitors for dynamic (CoA) requests from the specified server. See *Dynamic Service Management with RADIUS*.

Configuring Options that Apply to All RADIUS Servers

You can configure RADIUS options that apply to all RADIUS servers globally.

To configure RADIUS options globally:

1. Specify that you want to configure RADIUS options.

```
[edit access ]
user@host# edit radius-options
```

2. (Optional) Configure the rate at which RADIUS interim update requests are sent to the server.

```
[edit access radius-options]
user@host# set interim-rate interim-rate
```

3. (Optional) Configure the maximum allowed deviation from the configured update interval that the router sends interim accounting updates to the RADIUS server. The tolerance is relative to the configured update interval.

For example, if the tolerance is set to 60 seconds, then the router sends interim accounting updates no sooner than 30 seconds earlier than the configured update interval. When a subscriber logs in, the first interim accounting update may be sent up to 30 seconds early (on average 15 seconds early).

You configure the update interval with the *update-interval* statement at the [edit access profile *profile-name* accounting] hierarchy level.

```
[edit access radius-options]
user@host# set interim-update-tolerance seconds
```

4. (Optional) Configure the number of requests per second that the router can send to all configured RADIUS servers collectively. Limiting the flow of requests from the router to the RADIUS servers enables you to prevent the RADIUS servers from being flooded with requests.

```
[edit access radius-options]
user@host# set request-rate rate
```

5. (Optional) Configure the number of seconds that the router waits after a server has become unreachable before rechecking the connection. If the router reaches the server when the revert interval expires, the server is then used according to the order of the server list.

```
[edit access radius-options]
user@host# set revert-interval interval
```



NOTE: You can also configure the revert-interval in an access profile to override this global value. See ["Configuring Access Profile Options for Interactions with RADIUS Servers" on page 653](#).

6. (Optional) Configure the duration of a period during which unresponsive RADIUS authentication servers are not yet considered to be unreachable or down. You can vary the period depending on

whether you want to redirect authentication requests more quickly to another server or provide the unresponsive server more time to recover and respond.

See ["Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable"](#) on page 652 for more information.

```
[edit access radius-options]
user@host# set timeout-grace seconds
```

7. (Optional) Configure a NAS-Port value that is unique across all MX series routers in the network. You can configure a NAS-Port value that is unique within the router only, or unique across the different MX routers in the network.

See ["Enabling Unique NAS-Port Attributes \(RADIUS Attribute 5\) for Subscribers"](#) on page 694 for more information.

```
[edit access radius-options]
user@host# set unique-nas-port chassis-id chassis-id
user@host# set unique-nas-port chassis-id-width chassis-id-width
```

Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable

When a RADIUS authentication server fails to respond to any of the attempts for a given authentication request and times out, authd notes the time for reference, but it does not immediately mark the server as down (if other servers are available) or unreachable (if it is the only configured server). Instead, a configurable grace period timer starts at the reference time. The grace period is cleared if the server responds to a subsequent request before the period expires.

During the grace period, the server is not marked as down or unreachable. Each time the server times out for subsequent requests to that server, authd checks whether the grace period has expired. When the check determines that the grace period has expired and the server has still not responded to a request, the server is marked as unreachable or down.

Using a short grace period enables you to more quickly abandon an unresponsive server and direct authentication requests to other available servers. A long grace period gives a server more opportunities to respond and may avoid needlessly abandoning a resource. You might specify a longer grace period when you have only one or a small number of configured servers.

To configure the grace period during which an unresponsive RADIUS server is not marked as unreachable or down:

- Specify the duration of the grace period.

```
[edit access radius-options]
user@host# set timeout-grace seconds
```

Configuring Access Profile Options for Interactions with RADIUS Servers

You can use an access profile to specify options that the router uses when communicating with RADIUS authentication and accounting servers for subscriber access. This procedure describes options that are available only in access profiles. For options that are available at both the access profile and global level, see [RADIUS Servers and Parameters for Subscriber Access](#).

To configure RADIUS authentication and accounting server options:

1. Specify that you want to configure RADIUS options.

```
[edit access profile profile-name]
user@host# edit radius options
```

2. (Optional) Configure the format the router uses to identify the accounting session. The identifier can be in one of the following formats:
 - decimal—The default format. For example, 435264
 - description—In the format, `jnpr interface-specifier:subscriber-session-id`. For example, `jnpr fastEthernet 3/2.6:1010101010101`

```
[edit access profile profile-name radius options]
user@host# set accounting-session-id-format (decimal | description)
```

3. (Optional) Configure the delimiter character that the router inserts between values in RADIUS attribute 31 (Calling-Station-Id).

```
[edit access profile profile-name radius options]
user@host# set calling-station-id-delimiter "delimiter-character"
```

4. (Optional) Configure the information that the router includes in RADIUS attribute 31 (Calling-Station-Id).

See ["Configuring a Calling-Station-ID with Additional Options"](#) on page 660 for detailed information.

```
[edit access profile profile-name radius options]
user@host# set calling-station-id-format parameter
```

5. (Optional) Configure the router to use the optional behavior that inserts the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets, rather than sending the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets. This optional behavior requires that the value of the challenge must be 16 bytes; otherwise the statement is ignored and the challenge is sent as the CHAP-Challenge attribute.

```
[edit access profile profile-name radius options]
user@host# set chap-challenge-in-request-authenticator
```

6. (Optional) Configure the method the router uses to access RADIUS authentication and accounting servers when multiple servers are configured:
 - **direct**—The default method, in which there is no load balancing. The first server configured is the primary server; servers are accessed in order of configuration. If the primary server is unreachable, the router attempts to reach the second configured server, and so on.
 - **round-robin**—The method that provides load balancing by rotating router requests among the list of configured RADIUS servers. The server chosen for access is 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: When a RADIUS server in the round-robin list becomes unreachable, the next reachable server in the round-robin list is used for the current request. That same server is also used for the next request because it is at the top of the list of available servers. As a result, after a server failure, the server that is used takes up the load of two servers.

- To configure the method the router uses to access RADIUS accounting servers:

```
[edit access profile profile-name radius options]
user@host# set client-accounting-algorithm (direct | round-robin)
```

- To configure the method the router uses to access RADIUS authentication servers:

```
[edit access profile profile-name radius options]
user@host# set client-authentication-algorithm (direct | round-robin)
```

7. (Optional) Configure the router to use the optional behavior when a CoA operation is unable to apply a requested change to a client profile dynamic variable.

The optional behavior is that subscriber management does not apply any changes to client profile dynamic variables in the CoA request and then responds with a NACK. The default behavior is that subscriber management does not apply the incorrect update but does apply the other changes to the client profile dynamic variables, and then responds with an ACK message.

```
[edit access profile profile-name radius options]
user@host# set coa-dynamic-variable-validation
```

8. (Optional) Configure the router to use a physical port type of virtual to authenticate clients. The port type is passed in RADIUS attribute 61 (NAS-Port-Type). By default the router passes a port type of ethernet in RADIUS attribute 61.

```
[edit access profile profile-name radius options]
user@host# set ethernet-port-type-virtual
```



NOTE: This statement takes precedence over the `nas-port-type` statement if you include both in the same access profile.

9. (Optional) Specify the information that is excluded from the interface description that the router passes to RADIUS for inclusion in RADIUS attribute 87 (NAS-Port-ID). By default, the interface description includes adapter, channel, and subinterface information.

```
[edit access profile profile-name radius options]
user@host# set interface-description-format (exclude-adapter | exclude-channel | exclude-subinterface)
```

10. (Optional) For dual-stack PPP subscribers, include 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.

Optionally, configure a message that is included in the IPv4-Release-Control VSA (26–164) when it is sent to the RADIUS server

The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

```
[edit access profile profile-name radius options]  
user@host# set ip-address-change-notify message message
```

11. (Optional) 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.

```
[edit access profile profile-name radius options]  
user@host# set juniper-access-line-attributes
```

Starting in Junos OS Release 19.2R1, the `juniper-access-line-attributes` option replaces the `juniper-dsl-attributes` option. 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:

- Delete the `juniper-access-line-attributes` option from all access profiles that include it.
- Perform the software downgrade.
- Add the `juniper-dsl-attributes` option to the affected access profiles.

12. (Optional) Configure the value for the client RADIUS attribute 32 (NAS-Identifier), which is used for authentication and accounting requests.

```
[edit access profile profile-name radius options]
user@host# set nas-identifier identifier-value
```

13. (Optional) Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width of fields in the NAS-Port attribute, which specifies the physical port number of the NAS that is authenticating the user.

- For Ethernet subscribers:

```
[edit access profile profile-name radius options]
user@host# set nas-port-extended-format field width
```

- For ATM subscribers:

```
[edit access profile profile-name radius options]
user@host# set nas-port-extended-format atm field width
```

14. (Optional) Configure the delimiter character that the router inserts between values in RADIUS attribute 87 (NAS-Port-Id).

```
[edit access profile profile-name radius options]
user@host# set nas-port-id-delimiter delimiter-character
```

15. (Optional) Configure the optional information that the router includes in RADIUS attribute 87 (NAS-Port-Id). You can specify one or more options to appear in the default order. Alternatively, you can specify both the options and the order in which they appear. The orders are mutually exclusive and the configuration fails if you configure a NAS-Port-ID that includes values in both types of order.

See ["Configuring a NAS-Port-ID with Additional Options" on page 691](#) and ["Configuring the Order in Which Optional Values Appear in the NAS-Port-ID" on page 692](#) for detailed information.

```
[edit access profile profile-name radius options]
user@host# set nas-port-id-format optional-parameters
```


16. (Optional) Configure the port type that is included in RADIUS attribute 61 (NAS-Port-Type). This specifies the port type the router uses to authenticate subscribers.

```
[edit access profile profile-name radius options]  
user@host# set nas-port-type port-type
```



NOTE: This statement is ignored if you configure the ethernet-port-type-virtual in the same access profile.

17. (Optional) Configure the LAC to override the configured Calling-Station-ID format for the value sent in the L2TP Calling Number AVP 22. You can override the Calling-Station-ID format and configure the LAC to use the ACI, the ARI, or both the ACI and ARI that are received from the L2TP client in the PADR packet. You can also specify a delimiter to use between components of the AVP string and a fallback value to use when the configured override components are not received in the PADR packet.



NOTE: See [Override the Calling-Station-ID Format for the Calling Number AVP](#) for more information.

```
[edit access profile profile-name radius options]  
user@host# set override calling-station-id remote-circuit-id
```

18. (Optional) Override the value of the RADIUS NAS-IP-Address attribute (4) at the LNS with the value of the session's LAC endpoint IP address if it is present in the session database. If it is not present, the original attribute value is used.

```
[edit access profile profile-name radius options]  
user@host# set override nas-ip-address tunnel-client-gateway-address
```

19. (Optional) Override the value of the RADIUS NAS-Port attribute (5) at the LNS with the value from the session database if the LAC NAS port information was conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If it is not present, the original attribute value is used.

```
[edit access profile profile-name radius options]  
user@host# set override nas-port tunnel-client-nas-port
```

20. (Optional) Override the value of the RADIUS NAS-Port-Type attribute (61) at the LNS with the value from the session database if the LAC NAS port information was conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If it is not present, the original attribute value is used.

```
[edit access profile profile-name radius options]
user@host# set override nas-port-type tunnel-client-nas-port-type
```

21. (Optional) 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.



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.

```
[edit access profile profile-name radius options]
user@host# set remote-circuit-id-delimiter "delimiter"
```

22. (Optional) 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.

```
[edit access profile profile-name radius options]
user@host# set remote-circuit-id-fallback {configured-calling-station-id | default}
```

23. (Optional) Configure the format of the string that overrides the Calling-Station-ID format in the L2TP Calling Number AVP. You can specify the ACI, the ARI, or both the ACI and ARI.



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.

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

24. (Optional) Configure the number of seconds that the router waits after a server has become unreachable before making another attempt to reach the server. If the server is then reachable, it is used in accordance with the order of the server list.

```
[edit access profile profile-name radius options]  
user@host# set revert-interval interval
```



NOTE: You can also configure this option for all RADIUS servers. See [Configuring Options for RADIUS Servers](#).

25. (Optional) Configure whether newly authenticated subscriber can successfully log in when service activation failures related to configuration errors occur during authd processing of the activation request for the subscriber's address family. You can specify this behavior for services configured in dynamic profiles or in Extensible Subscriber Services Manager (ESSM) operation scripts:
- **optional-at-login**—Service activation is optional. Activation failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Service activation failures due to causes other than configuration errors cause network family activation to fail. The login attempt is terminated unless another address family is already active for the subscriber.
 - **required-at-login**—Service activation is required. Activation failure for any reason causes network family activation to fail. The login attempt is terminated unless another address family is already active for the subscriber.

```
[edit access profile profile-name radius options]  
user@host# set service-activation (dynamic-profile | extensible-service) (optional-at-login  
| required-at-login)
```

26. (Optional) Specify that RADIUS attribute 5 (NAS-Port) includes the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

```
[edit access profile profile-name radius options]  
user@host# set vlan-nas-port-stacked-format
```

Configuring a Calling-Station-ID with Additional Options

Use this section to configure an alternative value for the Calling-Station-ID (RADIUS IETF attribute 31) in an access profile on the MX Series router.

You can configure the Calling-Station-ID to include one or more of the following options, in any combination, at the [edit access profile *profile-name* radius options calling-station-id-format] hierarchy:

- Agent circuit identifier (*agent-circuit-id*)—Identifier of the subscriber's access node and the digital subscriber line (DSL) on the access node. The agent circuit identifier (ACI) string is stored in either the DHCP option 82 field of DHCP messages for DHCP traffic, or in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets for PPPoE traffic.
- Agent remote identifier (*agent-remote-id*)—Identifier of 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 either the DHCP option 82 field for DHCP traffic, or in the DSL Forum Agent-Remote-ID VSA [26-2] for PPPoE traffic.
- Interface description (*interface-description*)—Value of the interface.
- Interface text description (*interface-text-description*)—Text description of the interface. The interface text description is configured separately, using either the set interfaces *interface-name* description *description* statement or the set interfaces *interface-name* unit *unit-number* description *description* statement
- MAC address (*mac-address*)—MAC address of the source device for the subscriber.
- NAS identifier (*nas-identifier*)—Name of the NAS that originated the authentication or accounting request. NAS-Identifier is RADIUS IETF attribute 32.
- Stacked VLAN (*stacked-vlan*)—Stacked VLAN ID.
- VLAN (*vlan*)—VLAN ID.

If you configure the format of the Calling-Station-ID with more than one optional value, a hash character (#) is the default delimiter that the router uses as a separator between the concatenated values in the resulting Calling-Station-ID string. Optionally, you can configure an alternative delimiter character for the Calling-Station-ID to use. The following example shows the order of output when you configure multiple optional values:

```
nas-identifier#interface description#interface text description#agent-circuit-id#agent-remote-id#mac address#stacked vlan#vlan
```

To configure an access profile to provide optional information in the Calling-Station-ID:

1. Specify the access profile you want to configure.

```
[edit]
user@host# edit access profile profile-name
```

2. Specify that you want to configure RADIUS options.

```
[edit access profile profile-name]
user@host# edit radius options
```

3. Specify the nondefault character to use as the delimiter between the concatenated values in the Calling-Station-ID.

By default, subscriber management uses the hash character (#) as the delimiter in Calling-Station-ID strings that contain more than one optional value.

```
[edit access profile profile-name radius options]
user@host# set calling-station-id-delimiter delimiter-character
```

4. Configure the value for the NAS-Identifier (RADIUS attribute 32), which is used for authentication and accounting requests.

```
[edit access profile profile-name radius options]
user@host# set nas-identifier identifier-value
```

5. Specify that you want to configure the format of the Calling-Station-ID.

```
[edit access profile profile-name radius options]
user@host# edit calling-station-id-format
```

6. (Optional) Include the interface text description in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]
user@host# set interface-text-description
```

7. (Optional) Include the interface description value in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]
user@host# set interface-description
```

8. (Optional) Include the agent circuit identifier in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set agent-circuit-id
```

9. (Optional) Include the agent remote identifier in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set agent-remote-id
```

10. (Optional) Include the configured NAS identifier value in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set nas-identifier
```

11. (Optional) Include the stacked VLAN ID in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set stacked-vlan
```

12. (Optional) Include the VLAN ID in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set vlan
```

13. (Optional) Include the MAC address in the Calling-Station-ID.

```
[edit access profile profile-name radius options calling-station-id-format]  
user@host# set mac-address
```

Example: Calling-Station-ID with Additional Options in an Access Profile

The following example creates an access profile named `retailer01` that configures a Calling-Station-ID string that includes the NAS-Identifier (`fox`), interface description, agent circuit identifier, and agent remote identifier options.

```
[edit access profile retailer01 radius options]
nas-identifier "fox";
calling-station-id-delimiter "*";
calling-station-id format {
    nas-identifier;
    interface-description;
    agent-circuit-id;
    agent-remote-id;
}
```

The resulting Calling-Station-ID string is formatted as follows:

```
fox*ge-1/2/0.100:100*as007*ar921
```

where:

- The NAS-Identifier value is `fox`.
- The Calling-Station-ID delimiter character is `*` (asterisk).
- The interface description value is `ge-1/2/0.100:100`.
- The agent circuit identifier value is `as007`.
- The agent remote identifier value is `ar921`.

Consider an example where all options are configured, but no values are available for the Agent-Circuit-ID, the Agent-Remote-Id, or the stacked VLAN identifier. The other values are as follows:

- NAS identifier—`solarium`
- interface description—`ge-1/0/0.1073741824:101`
- interface text description—`example-interface`
- MAC address—`00:00:5E:00:53:00`
- VLAN identifier—`101`

These values result in the following Calling-Station-ID:

```
solarium#ge-1/0/0.1073741824:101#example-interface###00-00-5E-00-53-00##101
```

Filtering RADIUS Attributes and VSAs from RADIUS Messages

Standard attributes and vendor-specific attributes (VSAs) received in RADIUS messages take precedence over internally provisioned attribute values. Filtering attributes consists of choosing to *ignore* certain attributes when they are received in Access Accept packets and to *exclude* certain attributes from being sent to the RADIUS server. Ignoring attributes received from the RADIUS server enables your locally provisioned values to be used instead. Excluding attributes from being sent is useful, for example, for attributes that do not change for the lifetime of a subscriber. It enables you to reduce the packet size without loss of information.

You can specify standard RADIUS attributes and VSAs that the router or switch subsequently *ignores* when they are received in RADIUS Access-Accept messages. You can also specify attributes and VSAs that the router or switch *excludes* from specified RADIUS message types. Exclusion means that the router or switch does not include the attribute in specified messages that it sends to the RADIUS server.

Starting in Junos OS Release 18.1R1, you can configure the router or switch to ignore or exclude RADIUS standard attributes and VSAs by specifying the standard attribute number or the IANA-assigned vendor ID and the VSA number, respectively. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be ignored or excluded. The configuration has no effect if you configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to ignore only a subset of all attributes that can be received in Access-Accept messages.

To configure the attributes ignored or excluded by your router or switch:

1. Specify that you want to configure RADIUS in the access profile.

```
[edit access profile profile-name]
user@host# edit radius
```

2. Specify that you want to configure how RADIUS attributes are filtered.

```
[edit access profile profile-name radius]
user@host# edit attributes
```

3. (Optional) Specify one or more attributes you want your router or switch to ignore when the attributes are in Access-Accept messages.

- Legacy method: Specify dedicated option for attribute:

```
[edit access profile profile-name radius attributes]
user@host# set ignore attribute-name
```

- Flexible method: Specify standard attribute number or the IANA-assigned vendor ID and the VSA number:

```
[edit access profile profile-name radius attributes]
user@host# set ignore standard-attribute number
user@host# set ignore vendor-id id-number vendor-attribute vsa-number
```

4. (Optional) Configure an attribute that you want your router or switch to exclude from one or more specified RADIUS message types. You cannot configure a list of attributes, but you can specify a list of message types for each attribute.

- Legacy method: Specify dedicated option for attribute and message type:

```
[edit access profile profile-name radius attributes]
user@host# set exclude attribute-name [packet-type]
```

- Flexible method: Specify standard attribute number or the IANA-assigned vendor ID, the VSA number, and the message type:

```
[edit access profile profile-name radius attributes]
user@host# set exclude standard-attribute number packet-type [packet-type]
user@host# set exclude vendor-id id-number vendor-attribute vsa-number packet-type [packet-type]
```

The following example compares the legacy and flexible configuration methods to ignore the standard RADIUS attribute, Framed-IP-Netmask (9), and the Juniper Networks VSAs, Ingress-Policy-Name (26-10) and Egress-Policy-Name (26-11).

- Legacy method:

```
[edit access profile prof-ign radius attributes]
user@host# set ignore framed-ip-netmask input-filter output-filter
```

- Flexible method:

```
[edit access profile prof-ign radius attributes]
user@host# set ignore standard-attribute 9
user@host# set ignore vendor-id 4874 vendor-attribute [ 10 11 ]
```

The following example compares the legacy and flexible configuration methods to exclude the standard RADIUS attribute, Framed-IP-Netmask (9), and the Juniper Networks VSAs, Ingress-Policy-Name (26-10) and Egress-Policy-Name (26-11).

- Legacy method:

```
[edit access profile prof-exc radius attributes]
user@host# set exclude framed-ip-netmask accounting-stop
user@host# set exclude input-filter [ accounting-start accounting-stop ]
user@host# set exclude output-filter [ accounting-start accounting-stop ]
```

- Flexible method: Specify standard attribute number or the IANA-assigned vendor ID, the VSA number, and the message type:

```
[edit access profile prof-exc radius attributes]
user@host# set exclude standard-attribute 9 packet-type accounting-stop
user@host# set exclude vendor-id 4874 vendor-attribute 10 packet-type [ accounting-start accounting-stop ]
user@host# set exclude vendor-id 4874 vendor-attribute 11 packet-type [ accounting-start accounting-stop ]
```

What happens if you specify an attribute with both methods in the same profile? The effective configuration is the logical OR of the two methods. Consider the following example for the standard attribute, accounting-delay-time (41):

```
[edit access profile prof-3 radius attributes]
user@host# set exclude accounting-delay-time [ accounting-off accounting-on ]
user@host# set exclude standard-attribute 41 packet-type [ accounting-start accounting-stop ]
```

The result is that the attribute is excluded from all four message types: Accounting-Off, Accounting-On, Accounting-Start, and Accounting-Stop. The effect is the same as if either of the following configurations is used:

- [edit access profile prof-3 radius attributes]
user@host# **set exclude accounting-delay-time [accounting-off accounting-on accounting-start accounting-stop]**
- [edit access profile prof-3 radius attributes]
user@host# **set exclude standard-attribute 41 packet-type [accounting-off accounting-on accounting-start accounting-stop]**

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, you can configure the router or switch to ignore or exclude RADIUS standard attributes and VSAs by specifying the standard attribute number or the IANA-assigned vendor ID and the VSA number, respectively.

Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers

IN THIS SECTION

- [Interface Description Storage and Reporting Overview | 668](#)
- [Interface Description Storage and Reporting Configuration | 673](#)

Interface Description Storage and Reporting Overview

IN THIS SECTION

- [Interface Description Precedence | 669](#)
- [Example: Reporting Interface Descriptions on Non-Underlying Logical Interfaces | 670](#)

- [Reporting Interface Descriptions on Underlying Logical Interfaces | 671](#)
- [Example: PPP over an Underlying VLAN Demux Interface | 671](#)
- [Example: Reporting Interface Descriptions on Dynamic VLANs | 673](#)

You can configure Junos OS to store subscriber access interface descriptions and report the interface description through RADIUS. This capability enables you to uniquely identify subscribers on a particular logical or physical interface. When you enable storing of the interface descriptions, RADIUS requests include the interface description in VSA 26-63, if the subscriber's access interface has been configured with an interface description. All interface descriptions must be statically configured using the Junos OS CLI. Storing and reporting of interface descriptions is supported for DHCP, PPP, and authenticated dynamic VLANs, and applies to any client session that either authenticates or uses the RADIUS accounting service. The description can contain letters, numbers, and hyphens (-), and can be up to 64 characters long.

Interface Description Precedence

The interface description sent in the VSA depends on the configured interface. Two configuration models apply across topologies and protocols for subscriber management.

- Subscriber *logical interface* directly over a physical interface (non-underlying logical interfaces).
- Subscriber logical interface over an underlying logical interface and physical interface.

In both models, Junos OS selects the interface description to report based on order of precedence. Interfaces not configured with interface descriptions are excluded when selecting an interface by precedence. If no interface description is configured on any of the static interfaces in the subscriber interface hierarchy, VSA 26-63 is not sent in any of the RADIUS messages.



NOTE:

- For aggregated Ethernet physical interfaces, the interface description on the aggregated Ethernet interface, for example AE0 or AE1, serves as the physical interface description.
- If the subscriber's access is a combination of dynamic and static interfaces, Junos OS uses the description on the static interface.

Example: Reporting Interface Descriptions on Non-Underlying Logical Interfaces

This topic shows an example of subscriber access with non-underlying logical interfaces. In this case, the logical interface can be a VLAN or a VLAN demux interface. This example shows a DHCP subscriber logical interface over a VLAN without a demux interface. For non-underlying interfaces, Junos OS selects which interface description to report based on the following order of precedence:

1. Logical interface description
2. Physical interface description

Based on the order of precedence that Junos OS uses to select the interface description for non-underlying interfaces, Junos OS reports `subscriber_ifl_descr` as the interface description.

```
system {
  services {
    dhcp-local-server {
      group LSG1 {
        authentication {
          password $ABC123;
          username-include {
            user-prefix rich;
          }
        }
      }
      interface ge-1/0/0.100;
    }
  }
}

interfaces {
  ge-1/0/0 {
    description subscriber_ifd_descr;
    vlan-tagging;
    unit 100 {
      description subscriber_ifl_descr;
      vlan-id 100;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 198.51.100.20;
      }
    }
  }
}
```

Reporting Interface Descriptions on Underlying Logical Interfaces

Underlying logical interfaces can apply to both DHCP and PPP.

For DHCP, Junos OS selects which interface description to report based on the following order of precedence:

1. Underlying logical interface description
2. Underlying physical interface description



NOTE: For DHCP, Junos OS does not report the IP demux logical interface description.

For PPP over an underlying VLAN or VLAN demux interface, Junos OS selects which interface description to report based on the following order of precedence:

1. PPP interface description
2. Underlying VLAN without a demux interface or VLAN demux logical interface description
3. Underlying physical interface description

Example: PPP over an Underlying VLAN Demux Interface

The following example shows a PPP subscriber over an underlying VLAN demux interface. This configuration includes three possible interface descriptions. Based on the order of precedence that Junos OS uses to select the interface description for PPP, the interface description is reported as subscriber_ppp_ifl_descr_0.

```
interfaces {
  ge-1/0/0 {
    description subscriber_ifd_descr;
    hierarchical-scheduler maximum-hierarchy-levels 2;
    flexible-vlan-tagging;
  }
  demux0 {
    unit 0 {
      vlan-tags outer 1 inner 1;
      description subscriber_under_ifl_descr_1_1;
      demux-options {
        underlying-interface ge-1/0/0;
      }
      family pppoe {
```

```

        duplicate-protection;
    }
}
unit 1 {
    vlan-tags outer 1 inner 2;
    description subscriber_under_ifl_descr_1_2;
    demux-options {
        underlying-interface ge-1/0/0;
    }
    family pppoe {
        duplicate-protection;
    }
}
}
pp0 {
    unit 0 {
        description subscriber_ppp_ifl_descr_0;
        ppp-options {
            chap;
            pap;
        }
        pppoe-options {
            underlying-interface demux0.0;
            server;
        }
    }
    unit 1 {
        description subscriber_ppp_ifl_descr_1;
        ppp-options {
            chap;
            pap;
        }
        pppoe-options {
            underlying-interface demux0.1;
            server;
        }
    }
}
}
}

```

Example: Reporting Interface Descriptions on Dynamic VLANs

If you create dynamic VLANs with authentication, Junos OS reports the interface description on the physical interface. In the following example, dynamic VLANs created over the ge-1/2/0 interface are authenticated with an interface description of ge-1/2/0-bos-mktg-group.

```
ge-1/2/0 {
  description ge-1/2/0-bos-mktg-group;
  flexible-vlan-tagging;
  auto-configure {
    vlan-ranges {
      dynamic-profile vlan-prof {
        accept inet;
        ranges {
          any;
        }
      }
    }
    authentication {
      password $ABC123;
      username-include {
        user-prefix rich;
      }
    }
  }
}
```

Interface Description Storage and Reporting Configuration

To enable or disable storage and reporting of interface descriptions:

- Enable storing and reporting of interface descriptions.

```
[edit access]
```

```
user@host# set report-interface-descriptions
```


- Disable storing and reporting of interface descriptions per RADIUS message type.

```
[edit access profile profile-name radius attributes]
user@host# set exclude interface-description [ access-request | accounting-start | accounting-stop ]
```

RELATED DOCUMENTATION

[RADIUS Servers and Parameters for Subscriber Access](#) | 646

Session Options for Subscriber Access

IN THIS SECTION

- [Understanding Session Options for Subscriber Access](#) | 674
- [Subscriber Session Timeout Options](#) | 681
- [Limiting the Number of Active Sessions per Username and Access Profile](#) | 683
- [Configuring Username Modification for Subscriber Sessions](#) | 683
- [Removing Inactive Dynamic Subscriber VLANs](#) | 687

Session options enable you to specify several characteristics for DHCP, L2TP, and terminated PPP subscriber sessions. Session options are configured in access profiles that determine the parameters for subscriber access, authentication, authorization, and accounting.

Understanding Session Options for Subscriber Access

IN THIS SECTION

- [Subscriber Session Timeouts](#) | 675
- [Limits on Subscriber Sessions per Username and Access Profile](#) | 678
- [Benefits of Limiting Sessions for Usernames with the CLI](#) | 679

- Subscriber Username Modification | 680
- Benefits of Subscriber Username Modification | 681

You can use access profiles to configure several characteristics of the sessions that are created for DHCP, L2TP, and terminated PPP subscribers. You can place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both. You can limit subscriber sessions by username per access profile. You can also set parameters that modify a subscriber's username at login based on the subscriber's access profile.

Subscriber Session Timeouts

You can limit subscriber access by configuring a session timeout or an idle timeout. Use a session timeout to specify a fixed period of time that the subscriber is permitted to have access. Use an idle timeout to specify a maximum period of time that the subscriber can be idle. You can use these timeouts separately or together. By default, neither timeout is present.



NOTE: For all subscriber types other than *DHCP* (such as L2TP-tunneled and PPP-terminated subscribers), the session timeout value limits the subscriber session. For DHCP subscribers, the session timeout value is used to limit the lease when no other lease time configuration is present. The lease expires when the timeout value expires. If this value is not supplied by either the CLI or RADIUS, the DHCP lease does not expire.

The idle timeout is based on accounting statistics for the subscriber. The router determines subscriber inactivity 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.

Optionally, you can specify that only subscriber ingress traffic is monitored; egress traffic is ignored. This configuration 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 cannot detect that the peer is not up. In this situation, because by default the LAC monitors both ingress and egress traffic, it detects 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, the LAC can detect that the peer is inactive and then initiate logout.

When either timeout period expires, the non-DHCP subscribers are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout. DHCP subscribers are disconnected. The Acct-Terminate-Cause [RADIUS attribute 49] value includes a reason code of 5 for a session timeout and a code of 4 for an idle timeout.

You can configure these limitations to subscriber access on a per-subscriber basis by using the RADIUS attributes Session-Timeout [27] and Idle-Timeout [28]. RADIUS returns these attributes in Access-Accept messages in response to Access-Request messages from the access server. Starting in Junos OS Release 19.4R1, the Session-Timeout attribute [27] is supported in RADIUS CoA messages. This capability is useful, for example, when subscribers purchase Internet access for a specific period of time and must log out when the session expires.

When a CoA arrives with Session-Timeout, the timeout is counted from the time that the session activated. This has the following consequences:

- If the attribute value is greater than the current session uptime and between the minimum and maximum timeout values, the subscriber is logged out when that number of seconds has passed since session activation. For example, suppose the session activated at 12:00:00 and the CoA is received at 12:00:30 with a value of 120 seconds. The subscriber is logged out at 12:02:00.

Another way to look at this with the same values is that the current session uptime is 30 seconds and the attribute value is 120 seconds. The subscriber is logged out when 90 more seconds have passed.

- If the attribute value is greater than the current session uptime but less than the minimum timeout value of 60 seconds, then the subscriber is logged out when the uptime reaches 60 seconds.
- If the attribute value is greater than the current session uptime but more than the maximum timeout value of 31,622,400 seconds, then the subscriber is logged out when the uptime reaches 31,622,400 seconds.
- If the attribute value is less than the current session uptime, the session timeout is not applied. AAA replies to the CoA message with a NAK. For example, the session is unaffected if the Session-Timeout is 60 seconds, but the uptime is 100 seconds.

Applying a session timeout according to the rules above also depends on whether all other aspects of the CoA are successful. For example, if the CoA includes a service activation and that service activation fails, then the session timeout is not applied. AAA replies to the CoA message with a NAK.



NOTE: If the Session-Timeout value is 0, then any existing session timeout for that session is cancelled.

Service providers often choose to apply the same limitations to large numbers of subscribers. You can reduce the RADIUS provisioning effort for this scenario by defining the limitations for subscribers in an access profile on a per-routing-instance basis. If you do so, RADIUS attributes subsequently returned for a particular subscriber logged in with the profile override the per-routing-instance values.



BEST PRACTICE: We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is based only on time

and not user activity, 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.



BEST PRACTICE: We recommend that you do not configure an idle timeout for DHCP subscribers. When the timeout expires with no activity and the connection is terminated, the protocol has no means to inform the client. Consequently, these subscribers are forced to reboot their CPE device the next time they attempt to access the Internet.

Contrast the behavior when an idle timeout is configured for PPP subscribers. In this case, timeout expiration causes PPP to terminate the link with the peer. Depending on the CPE device, this termination enables the peer to automatically retry the connection either on demand or immediately. In either case, no subscriber intervention is required.

The available range for setting a timeout is the same whether you configure it in the CLI or through the RADIUS attributes:

- Session timeouts can be set for 1 minute through 527,040 minutes in the CLI and the corresponding number of seconds (60 through 31,622,400) in the Session-Timeout attribute [27].
- Idle timeouts can be set for 10 minutes through 1440 minutes in the CLI and the corresponding number of seconds (600 through 86,400) in the Idle-Timeout attribute [28].

The router interprets the values in the attributes to conform to the supported ranges. For example, for Session-Timeout [27]:

- A value of zero is treated as no timeout.
- A value in the range 1 through 59 is raised to 60 seconds.
- A value that exceeds 31,622,400 is reduced to 31,622,400 seconds.

For Idle-Timeout [28]:

- A value of zero is treated as no timeout.
- A value in the range 1 through 599 is raised to 600 seconds.
- A value that exceeds 86,400 is reduced to 86,400 seconds.

In configurations using dynamically created subscriber VLANs, the idle timeout also deletes the inactive subscriber VLANs when the inactivity threshold has been reached. In addition to deleting inactive dynamic subscriber VLANs, the idle timeout also removes dynamic VLANs when no client sessions were ever created (for example, in the event no client sessions are created on the dynamic VLAN or following

the occurrence of an error during session creation or client authentication where no client sessions are created on the dynamic VLAN).

Session and idle timeouts for deleting dynamic subscriber VLANs are useful only in very limited use cases; typically neither timeout is configured for this purpose.

A possible circumstance when they might be useful is when the dynamic VLANs have no upper layer protocol that helps determine when the VLAN is removed with the `remove-when-no-subscribers` statement; for example, when the VLAN is supporting IP over Ethernet without DHCP in a business access model with fixed addresses. However, business access is generally a higher-tier service than residential access and as such typically is not subject to timeouts due to inactivity as might be desired for residential subscribers.

An idle timeout might be appropriate in certain Layer 2 wholesale situations, where the connection can be regenerated when any packet is received from the CPE.

When using the idle timeout for dynamic VLAN removal, keep the following in mind:

- The idle timeout period begins after a dynamic subscriber VLAN interface is created or traffic activity stops on a dynamic subscriber VLAN interface.
- If a new client session is created or a client session is reactivated successfully, the client idle timeout resets.
- The removal of inactive subscriber VLANs functions only with VLANs that have been authenticated.

Limits on Subscriber Sessions per Username and Access Profile

Legitimate subscribers might share their login credentials with unauthorized persons, expending service provider resources without benefit to the provider. Starting in Junos OS Release 18.4R1, you can control or prevent the sharing of login credentials by limiting the number of active subscriber sessions that are allowed for a specific username associated with an access profile. You can also achieve this control with RADIUS, but configuring the limit locally on the BNG eliminates dependency on an external server.

When you configure a limit, active sessions for the username/access profile combination are tracked. The number of tracked sessions is checked when `authd` receives a new session login request. If the number of tracked session matches the limit, the new login attempt is rejected and counted as a blocked request.

When `authd` receives a logout or client termination request for a session, the tracked-sessions count is decremented for that username/access profile entry. If this continues until there are no active sessions for the combination, the entry is removed from the session limit table. All associated username/access profile entries are removed from the table if you delete the access profile or the session-limit from your configuration.

The total number of sessions for a username can exceed the configured limit for a particular access profile, because the same username can be used with multiple access profiles.



NOTE: For stacked subscriber sessions such as PPP with autoconfigured VLANs, both usernames in the stack are used for authentication and consequently both are counted against the session limit.

The configured limit applies to existing active subscribers, but existing sessions are not torn down if number of active sessions exceeds the limit for a subscriber with that username and access profile combination.

Consider a situation where five sessions are currently active for a given username/access profile combination when you configure a limit of two.

1. The active sessions count is recorded as five in the session limit table entry for the combination.
2. A new subscriber with the same username and access profile tries to log in. The attempt is blocked because the limit of two sessions is already exceeded (five > two).
3. An existing subscriber logs out, decrementing the active sessions count to four.
4. A new subscriber with the same username and access profile tries to log in. The attempt is blocked because the limit of two sessions is still exceeded (four > two).
5. Three existing subscribers log out, reducing the active sessions count to one.
6. A new subscriber with the same username and access profile tries to log in. The attempt is allowed because the limit of two sessions has not yet been reached (one < two).

The session limit design prevents a denial-of-service event where a malicious user makes multiple login attempts with the correct username and access profile, but the wrong password. The numerous login attempts might exceed the configured session limit, but this does not occur because the tracked-sessions count is incremented only when the subscriber session state transitions to the active state, which the malicious logins do not achieve.

Benefits of Limiting Sessions for Usernames with the CLI

- Enables you to limit the number of sessions locally on the router, rather than being dependent on an external RADIUS server to provide the limit.
- Prevents some denial-of-service attacks based on multiple logins.



Subscriber Username Modification

For Layer 2 wholesale applications, some network service providers employ username modification to direct subscribers to the appropriate retail enterprise network. This modification is also called username *stripping*, because some of the characters in the username are stripped away and discarded. The remainder of the string becomes the new, modified username. The modified username is used by an external AAA server for session authentication and accounting. The modification parameters are applied according to a subscriber access profile that also determines the subscriber and session context; that is, the logical system:routing instance (LS:RI) used by the subscriber. Only the default (primary) logical system is supported. Because the wholesaler differentiates between multiple retailers by placing each in a different LS:RI, the usernames are appropriately modified for each retailer.

You can select up to eight characters as delimiters to mark the boundary between the discarded and retained portions of the original username; there is no default delimiter. The portion of the name to the right of the selected delimiter is discarded along with the delimiter. By configuring multiple delimiters, a given username structure can result in different modified usernames. You can configure the direction in which the original name is parsed to determine which delimiter marks the boundary. By default, the parse direction is from left to right.



Consider the following examples:

- You specify one delimiter, @. The username is user1@example.com. In this case, the parse direction does not matter. In either case, the single delimiter is found and example.com is discarded. The modified username is user1.

parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1



8043376

- You specify one delimiter, @. The username is user1@test@example.com. In this case, the parse direction results in different usernames.
 - Parse direction is left-to-right—The left-most @ is identified as the delimiter and test@example.com is discarded. The modified username is user1.
 - Parse direction is right-to-left—The right-most @ is identified as the delimiter and example.com is discarded. The modified username is user1@test.

parse direction	identify delimiter	modified username
left-to-right	user1@ test @example.com 	user1
right-to-left	user1@test@ example .com 	user1@test

8043377

- You specify two delimiters, @ and /. The username is user1@bldg1/example.com. The parse direction results in different usernames.
- Parse direction is left-to-right—The @ is identified as the delimiter and bldg1/example.com is discarded. The modified username is user1.
- Parse direction is right-to-left—The / is identified as the delimiter and example.com is discarded. The modified username is user1@bldg1.

parse direction	identify delimiter	modified username
left-to-right	user1@bldg1/ example .com 	user1
right-to-left	user1@bldg1/ example .com 	user1@bldg1

8043378

You can configure a subscriber access profile so that a portion of each subscriber login string is stripped and subsequently used as a modified username by an external AAA server for session authentication and accounting. The modified username appears, for example, in RADIUS Access-Request, Acct-Start, and Acct-Stop messages, as well as RADIUS-initiated disconnect requests and change of authorization (CoA) requests.

Benefits of Subscriber Username Modification

- Enables Layer 2 wholesale network service providers to easily direct subscribers to the appropriate retail enterprise network.

SEE ALSO

[RADIUS IETF Attributes Supported by the AAA Service Framework](#) | 4

Subscriber Session Timeout Options

Subscriber session timeout options enable you to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both. The subscriber session options

apply to both L2TP-tunneled and PPP-terminated subscriber sessions. For DHCP subscribers, the session timeout limits the DHCP lease time.



NOTE: To configure the timeout attributes in RADIUS, refer to the documentation for your RADIUS server.

To configure limitations on subscriber sessions, configure the session options in the client profile that applies to the subscriber:

- Terminate the subscriber when the configured session timeout expires, regardless of activity.

```
[edit access profile profile-name session-options]
user@host# set client-session-timeout minutes
```

- Terminate the subscriber when there is no ingress or egress data traffic for the duration of the configured idle timeout.

```
[edit access profile profile-name session-options]
user@host# set client-idle-timeout minutes
```

- Terminate the subscriber when there is no ingress data traffic for the duration of the configured idle timeout; ignore egress traffic.

```
[edit access profile profile-name session-options]
user@host# set client-idle-timeout minutes
user@host# set client-idle-timeout-ingress-only
```

For example, to configure session timeout options in the `acc-prof` client profile, specifying an idle timeout of 15 minutes, that only ingress traffic is monitored, and that the session times out after 120 minutes:

```
[edit]
access {
  profile {
    acc-prof {
      session-options {
        client-idle-timeout 15;
        client-idle-timeout-ingress-only;
        client-session-timeout 120;
      }
    }
  }
}
```

```

    }
  }
}

```

Limiting the Number of Active Sessions per Username and Access Profile

You can control the degree to which legitimate subscribers can share their login credentials by limiting the number of active subscriber sessions that are allowed for a specific username associated with an access profile.

To limit the number of active sessions per username and access profile:

-

```

[edit access profile profile-name]
user@host# set session-limit-per-username number

```

For example, to set the maximum number of active sessions per username to five for the access profile `isp-weg-4`:

```

[edit access profile isp-weg-4]
user@host# set session-limit-per-username 5

```

You can use the `show network-access aaa statistics session-limit-per-username` command to view statistics for active sessions and blocked requests.

You can use the `clear network-access aaa statistics session-limit-per-username username` command as an aid to debugging by clearing the blocked request statistics for any of the following cases:

- For all usernames across all access profiles.
- For a specific username across all access profiles.
- For a specific username in a specific access profile.
- For all usernames in a specific access profile.

Configuring Username Modification for Subscriber Sessions

You can use subscriber session options to set parameters that modify a subscriber's username at login based on the subscriber's access profile. This modification is also called username *stripping*, because some of the characters in the username are stripped away and discarded. The remainder of the string becomes the new, modified username. The modified username is used by an external AAA server for session authentication and accounting. This capability can be useful, for example, in Layer 2 wholesale

implementations, where the network service providers employ username modification to direct subscribers to the appropriate retail enterprise network.

The modification parameters are applied according to a subscriber access profile that also determines the subscriber and session context; that is, the logical system:routing instance (LS:RI) used by the subscriber. Only the default (primary) logical system is supported. Because the wholesaler differentiates between multiple retailers by placing each in a different LS:RI, the usernames are appropriately modified for each retailer.

You can select up to eight characters as delimiters to mark the boundary between the discarded and retained portions of the original username; there is no default delimiter. The portion of the name to the right of the selected delimiter is discarded along with the delimiter. By configuring multiple delimiters, a given username structure can result in different modified usernames. You can configure the direction in which the original name is parsed to determine which delimiter marks the boundary. By default, the parse direction is from left to right.

To configure username modification:

1. Define a profile consisting of a set of AAA options for authorizing and configuring a subscriber or set of subscribers with a subscriber access profile.

- a. Specify the name of the subscriber access profile that includes the username stripping configuration.

```
[edit access aaa-options aaa-options-name]
user@host# access-profile profile-name
```

- b. (Optional) Specify the logical-system:routing-instance (LS:RI) that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting. For example, this may correspond to the LS:RI for a retail ISP that provides services to the subscriber.

```
[edit access aaa-options aaa-options-name]
user@host# aaa-context aaa-context-name
```

- c. (Optional) Specify the logical-system:routing-instance (LS:RI) in which the subscriber interface is placed. For example, this may correspond to the LAC-facing interface on the LNS that is accessed by all requests from a subscriber residence.

```
[edit access aaa-options aaa-options-name]
user@host# subscriber-context subscriber-context-name
```

2. Configure the session options in the access profile that specify how usernames are stripped.

- a. Specify one or more delimiters to mark the boundary between the discarded and retained portions of the original username.

```
[edit access profile profile-name session-options strip-user-name]
user@host# set delimiter [ delimiter ]
```

- b. (Optional) Specify the direction in which the original username is examined to find a delimiter. The default direction is left-to-right.

```
[edit access profile profile-name session-options strip-user-name]
user@host# set parse-direction (left-to-right | right-to-left)
```

3. (Optional) Specify that the AAA options are on a per-interface basis when dynamic subscribers are authenticated.

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

4. (Optional) Specify that the AAA options are part of the PPP options in a group profile that applies to tunneled PPP subscribers at the LNS.

```
[edit access group-profile profile-name ppp]
user@host# set ppp-options aaa-options aaa-options-name
```



In the following example, the AAA options profile, `aaa1`, specifies a subscriber access profile, `entA`, for subscribers in the default logical system and routing instance 1. The access profile, `entA`, specifies that usernames are examined from left to right until the delimiter, `@`, is found. The AAA options profile is applied to tunneled PPP subscribers that belong to the group profile, `FD1`.

```
[edit access aaa-options aaa1]
user@host# access-profile entA
user@host# aaa-context default:1

[edit access profile entA session-options strip-user-name]
user@host# set delimiter @
user@host# set parse-direction left-to-right
```

```
[edit access group-profile FD1 ppp]
user@host# set ppp-options aaa-options aaa1
```

Given that configuration, suppose a subscriber attempts to log in with the username, user1@example.com. When this name is examined, the delimiter and the string example.com are discarded, leaving a modified username of user1. Note that the result is the same if the parse direction is set to examine the name from right to left, because only one delimiter is defined and only one is present in the original username.

parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1



8043376

Now suppose the subscriber logs in with the username, user1@test@example.com. For a username like this, the parsing direction makes a difference in the modified username. The configuration determines that the first instance of the delimiter @ is found first, because the name is parsed from left to right. This delimiter and the string test@example.com are discarded, leaving user1 as the modified username.

What happens when the configuration sets a different parsing direction?

```
[edit access profile entA session-options strip-user-name]
user@host# set delimiter @
user@host# set parse-direction right-to-left
```

In this case, for the username user1@test@example.com, the second instance of the delimiter is identified and it is discarded with the string @example.com. The modified username is user1@test.



parse direction	identify delimiter	modified username
left-to-right		user1
right-to-left		user1@test

8043377

You can achieve the same results of different modified usernames based on parse direction by configuring more than one delimiter as in the following configuration, where you specify two delimiters, @ and /.

```
[edit access profile entA session-options strip-user-name]
user@host# set delimiter [@ /]
user@host# set parse-direction left-to-right
```

For the username user1@bldg1/example.com, parsing left to right identifies the @ delimiter first and the modified username is user1. Parsing right to left instead, identifies the / delimiter first and strips it away with the string example.com, leaving a modified username of user1@bldg1.

parse direction	identify delimiter	modified username
left-to-right	user1@bldg1/example.com 	user1
right-to-left	user1@bldg1/example.com 	user1@bldg1

8043378


SEE ALSO

- [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile](#)
- [Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface](#)

Removing Inactive Dynamic Subscriber VLANs

Subscriber session timeouts enable you to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both. In configurations using dynamically created subscriber VLANs, the idle timeout also:

- Deletes the inactive subscriber VLANs when the inactivity threshold has been reached.
- Removes dynamic VLANs when no client sessions were ever created (for example, in the event no client sessions are created on the dynamic VLAN or following the occurrence of an error during session creation or client authentication where no client sessions are created on the dynamic VLAN).



NOTE: Session timeouts are typically not used for deleting dynamic subscriber VLANs. The timeout might be useful only in very limited use cases. One case might be when the dynamic VLANs have no upper layer protocol that helps determine when the VLAN is removed with the `remove-when-no-subscribers` statement; for example, when the VLAN is

supporting IP over Ethernet without DHCP in a business access model with fixed addresses.



NOTE: To configure the idle timeout attribute in RADIUS, refer to the documentation for your RADIUS server.

To remove inactive dynamic subscriber VLANs:

- 1. Edit session options for the router access profile.

```
[edit]
user@host# edit access profile profile-name session-options
```

- 2. Configure the maximum period a subscriber session can remain idle.

```
[edit access profile profile-name session-options]
user@host# set client-idle-timeout minutes
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.4R1	Starting in Junos OS Release 19.4R1, the Session-Timeout attribute [27] is supported in RADIUS CoA messages.
18.4R1	Starting in Junos OS Release 18.4R1, you can control or prevent the sharing of login credentials by limiting the number of active subscriber sessions that are allowed for a specific username associated with an access profile.

RELATED DOCUMENTATION

| [RADIUS Servers and Parameters for Subscriber Access](#) | 646

RADIUS NAS Port Attributes and Options

IN THIS SECTION

- [Manual Configuration of the NAS-Port-ID RADIUS Attribute | 689](#)
- [Configuring a NAS-Port-ID with Additional Options | 691](#)
- [Configuring the Order in Which Optional Values Appear in the NAS-Port-ID | 692](#)
- [Enabling Unique NAS-Port Attributes \(RADIUS Attribute 5\) for Subscribers | 694](#)
- [RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview | 695](#)
- [Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 697](#)
- [Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 698](#)
- [Manual Configuration of the NAS-Port-Type RADIUS Attribute | 699](#)
- [Configuring the RADIUS NAS-Port-Type per Physical Interface | 702](#)
- [Configuring the RADIUS NAS-Port-Type per VLAN | 703](#)
- [Configuring the RADIUS NAS-Port-Type per Stacked VLAN | 705](#)
- [Configuring the RADIUS NAS-Port Extended Format per Physical Interface | 707](#)
- [Configuring the RADIUS NAS-Port Extended Format per VLAN | 708](#)
- [Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN | 710](#)
- [Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces | 712](#)

Manual Configuration of the NAS-Port-ID RADIUS Attribute

Subscriber management uses the NAS-Port-ID (RADIUS attribute 87) to provide an interface description that identifies the physical interface that is used to authenticate subscribers. The NAS-Port-ID is included in RADIUS Access-Request, Acct-Start, Acct-Stop, Acct-On, and Acct-Off messages.

You can configure access profiles to specify additional information in the NAS-Port-ID. The additional information can be any combination of the interface description (the default value), the Agent Circuit ID, the Agent Remote ID, and the NAS identifier. You can also specify an optional delimiter character, which separates the values in a NAS-Port-ID. The default delimiter character is the hash character (#).

The NAS-Port-ID for nonchannelized interfaces consists of an interface-description string with one of the following formats:

- Default format:

interface-type slot/adapter/port.subinterface[:svlan-vlan]

For example, ge-1/2/0.100:100.

- Format when you use a demux VLAN as the underlying logical interface:

interface-type slot/adapter/port.demux0.subinterface[:svlan-vlan]

For example, ge-1/2/0.demux0.100:100-100

- Format when you use a demux VLAN as the underlying logical interface for an aggregated Ethernet interface:

ae interface-number.demux0.subinterface[:svlan-vlan]

For example, ae1.demux0.101:100-101

Starting in Junos OS Release 17.3R1, a logical port number is added to the default format for only channelized interfaces. For channelized interfaces, the default format for a NAS-Port-ID consists of the following interface-description string:

interface-type slot/adapter/logical-port-number.subinterface[:svlan-vlan]

For example, xe-0/1/143.4-5.6.

You can optionally configure the interface description format in an access profile to exclude the adapter, channel, or subinterface information.

You might optionally configure an access profile that specifies that the NAS-Port-ID includes the NAS identifier, the Agent Circuit ID, and the Agent Remote ID, in addition to the default interface description. For this configuration, the NAS-Port-ID consists of the following string:

nas-identifier#interface-description#agent-circuit-id#agent-remote-id

For example:

retailer25#ge-1/2/0.100:100#ACI 12/1/22/1230:1.1.23#ARI 55/2/23.9999:10.11.1923



NOTE: The NAS-Port-ID displays the configured values in the following order (where # is the delimiter):

nas-identifier#interface-description#agent-circuit-id#agent-remote-id

Configuring a NAS-Port-ID with Additional Options

The NAS-Port-ID (RADIUS attribute 87) identifies the physical interface that subscriber management uses to authenticate subscribers. By default, the NAS-Port-ID includes the interface-description value that describes the physical interface. You can include the following optional values in the NAS-Port-ID:

- agent-circuit-id
- agent-remote-id
- interface-description
- interface-text-description
- nas-identifier
- postpend-vlan-tags



NOTE: If you specify any optional values, the default interface-description value is no longer automatically included. You must explicitly specify the interface-description value if you want it to appear in the NAS-Port-ID.

When you specify optional values, the router arranges the values in the following default order, where the # character is the default delimiter:

```
nas-identifier # interface-description # interface-text-description # agent-circuit-id # agent-remote-id #
postpend-vlan-tags
```

You can use the order option to configure the explicit order in which the specified optional values appear in the NAS-Port-ID string.

To configure optional values in the NAS-Port-ID string:

1. Specify the access profile you want to configure.

```
[edit]
user@host# edit access profile retailer25
```

2. Specify that you want to configure RADIUS options.

```
[edit access profile retailer25]
user@host# edit radius options
```

3. Specify the character to use as the delimiter between the different attribute values in the NAS-Port-ID. By default, subscriber management uses the hash character (#).

```
[edit access profile retailer25 radius options]
user@host# set nas-port-delimiter %
```

4. Specify that you want to configure the format of the NAS-Port-ID.

```
[edit access profile retailer25 radius options]
user@host# edit nas-port-id-format
```

5. (Optional) Specify the optional values you want to include in the NAS-Port-ID string. The optional values appear in the default order.

```
[edit access profile retailer25 radius options nas-port-id-format]
user@host# set interface-description nas-identifier agent-remote-id agent-circuit id
```

6. (Optional) To specify an explicit non-default order in which the optional values appear in the NAS-Port-ID string, include the order option before each optional value. Specify the values in the order you want them to appear.

See ["Configuring the Order in Which Optional Values Appear in the NAS-Port-ID" on page 692](#).

Configuring the Order in Which Optional Values Appear in the NAS-Port-ID

In addition to specifying the values that you want to include in the NAS-Port-ID, you can use the order option to specify the explicit order in which you want the values to appear.

By default, the router arranges the specified values in the following order, where the # character is the delimiter:

```
nas-identifier # interface-description # interface-text-description # agent-circuit-id # agent-remote-id # postpend-vlan-tags
```



NOTE: The default order and the customized order are mutually exclusive. The configuration fails if you try to specify both.

To configure the specific order in which you want the optional values to appear in the NAS-Port-ID:

1. Specify that you want to configure the format of the NAS-Port-ID.

```
[edit access profile retailer25 radius options]
user@host# edit nas-port-id-format
```

2. Include the `order` option before each optional value that you want to include in the NAS-Port-ID. Specify the optional values in the order in which you want them to appear.

```
[edit access profile retailer25 radius options nas-port-id-format]
user@host# set order interface-description order nas-identifier order agent-remote-id order interface-
text-description
```

This configuration configures the following NAS-Port-ID string, where the % character is the delimiter:

```
interface-description % nas-identifier % agent-remote-id % interface-text-description
```

3. (Optional) To add an optional value to an existing NAS-Port-ID string:

Use the `order` option and the name of the optional value to add the new value to the existing NAS-Port-ID. The new value is added at the end of the string. For example:

```
[edit access profile retailer25 radius options nas-port-id-format]
user@host# set order agent-circuit-id
```

This configuration modifies the example in the previous step by adding the `agent-circuit-id` to the end of the NAS-Port-ID string:

```
interface-description % nas-identifier % agent-remote-id % interface-text-description % agent-circuit-id
```



NOTE: If you attempt to add an optional value that already exists in the NAS-Port-ID string, the new specification is ignored and the existing value remains in the order in which it was originally configured.

If you want to modify the existing order, delete the existing specification and define the new order.

Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers

Typically, the router derives the RADIUS NAS-Port attribute (attribute 5) value from a subscriber's physical port, as shown in the following list.

- Subscribers over Ethernet interfaces—combination of slot/adaptor/port/SVLAN ID/VLAN ID
- Subscribers over ATM interfaces—combination of slot/adaptor/port/VPI/VCI

However, in some customer environments, a NAS-Port attribute that is based on the physical port might not be unique, and multiple subscribers might have the same NAS-Port value. To avoid the duplicate use of a NAS-Port attribute, you can configure the router to provide unique NAS-Port attributes. The unique NAS-Port attribute consists of 32 bits (the most significant bit [MSB] is always 0), which make up two parts— a unique number that the router internally generates, and an optional unique chassis ID that you specify.

If you create the NAS-Port value based on the internally generated number only, the resulting NAS-Port value is unique within the router only. If your implementation requires NAS-Port values to be unique across all MX series routers in the network, you must also configure the unique chassis ID.

Uniqueness across all routers—To configure a NAS-Port attribute that is unique across all routers in the network, you use the following procedure:

- Configure the chassis ID width (1–7 bits)—You must use the same width for all routers in the network.
- Configure the chassis ID—You must ensure that you configure a unique ID for each router.
- The router uses the remainder of the 31 bits (minus the MSB and the number of bits used for the chassis ID width) for the internally generated number.

Uniqueness within the local router—To configure a NAS-Port attribute that is unique within the local router only, you use the following procedure:

- Do not configure the chassis ID width or chassis ID.
- The router uses all 31 bits for the internally generated number. The resulting NAS-Port attribute is unique only within the router and cannot be guaranteed to be unique for any other routers in the network.

To configure unique NAS-Port attribute values for subscribers:



NOTE: Before configuring the unique NAS-Port attribute, ensure that neither the `nas-port-extended-format` statement or the `vlan-nas-port-stacked-format` statement is configured at the [edit access profile *profile-name* radius options] hierarchy level. Otherwise, the commit operation will fail.

1. Specify that you want to configure RADIUS options at the [edit access] hierarchy level.

```
[edit access]
user@host# edit radius-options
```

2. Specify that you want to enable unique NAS-Port attribute support.

```
[edit access radius-options]
user@host# edit unique-nas-port
```



NOTE: This step configures the router to generate a unique number, which creates a NAS-Port value that is unique within the router.

3. (Optional) If you want to provide NAS-Port values that are unique across all MX series routers in the network, complete the following additional steps.
 - Specify the number of bits used in the chassis ID portion of the NAS-Port attribute. You can specify 1-7 bits. You must use the same chassis ID width for all routers across the network.

```
[edit access radius-options unique-nas-port]
user@host# set chassis-id-width chassis-id-width
```

- Specify the value you want to use for chassis ID portion of the NAS-Port attribute. The chassis ID can be in the range from 0-127 bits. You must configure a unique chassis ID for each MX router in the network.

```
[edit access radius-options unique-nas-port]
user@host# set chassis-id chassid-id
```

RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview

IN THIS SECTION

- [NAS-Port-Type RADIUS Attribute | 696](#)
- [NAS-Port RADIUS Attribute | 696](#)

- [NAS-Port Options Configuration and Subscriber Network Access Models | 696](#)
- [NAS-Port Options Definition | 697](#)

On MX Series routers with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces, you can configure the NAS-Port-Type (61) RADIUS IETF attribute, and an extended format for the NAS-Port (5) RADIUS IETF attribute, on a per-interface, per-VLAN, or per-stacked VLAN basis. The router passes the NAS-Port and NAS-Port-Type attributes to the RADIUS server during the authentication, authorization, and accounting (AAA) process.

This overview covers the following topics:

NAS-Port-Type RADIUS Attribute

The NAS-Port-Type attribute specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber. When you use the `nas-port-type` statement to configure the NAS-Port-Type, you can specify one of several predefined port types, or a user-defined port type value in the range 0 through 65535.

NAS-Port RADIUS Attribute

The NAS-Port attribute specifies the physical port number of the NAS that is authenticating the user, and is formed by a combination of the physical port's slot number, port number, adapter number, VLAN ID, and S-VLAN ID. The NAS-Port extended format, which you configure with the `nas-port-extended-format` statement, specifies the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, VLAN, and S-VLAN.

To include stacked VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format, use the `stacked` option as part of the `nas-port-extended-format` statement. If you do not configure the `stacked` option, stacked VLAN IDs are not included in the extended format.

NAS-Port Options Configuration and Subscriber Network Access Models

Configuring the NAS-Port-Type and the extended format for NAS-Port on a per-VLAN, per-stacked VLAN, or per-physical interface basis is useful in network configurations that use the following subscriber access models:

- **1:1 access model (per-VLAN basis)**—In a 1:1 access model, dedicated customer VLANs (C-VLANs) provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.

- N:1 access model (per-S-VLAN basis)—In an N:1 access model, service VLANs are dedicated to a particular service, such as video, voice, or data, instead of to a particular subscriber. Because a service VLAN is typically shared by many subscribers within the same household or in different households, the N:1 access model provides a many-to-one correspondence between individual subscribers and the VLAN encapsulation.
- 1:1 or N:1 access model (per-physical interface basis)—You can configure the NAS-Port-Type and NAS-Port format on a per-physical interface basis for both the 1:1 access model and the N:1 access model.

NAS-Port Options Definition

As an alternative to globally configuring the NAS-Port-Type and NAS-Port extended format in an access profile, you can configure these attributes on a per-interface, per-VLAN, or per-stacked VLAN basis. To do so, you must create a *NAS-Port options definition*, which includes some or all of the following components:

- NAS-Port-Type value—Specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.
- NAS-Port extended format—Configures the number of bits (bit width) for each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN. Optionally, you can also use the stacked option as part of the `nas-port-extended-format` statement to include S-VLAN IDs, in addition to VLAN IDs, in the extended format. If you do not configure the stacked option, stacked VLAN IDs are not included in the extended format.
- VLAN ranges or S-VLAN ranges—Defines the VLAN range of subscribers or stacked VLAN range of subscribers to which each NAS-Port options definition applies.

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

The following guidelines apply when you configure the NAS-Port-Type attribute and the extended format for the NAS-Port attribute on a per-VLAN, per-stacked VLAN, or per-physical interface basis:

- You can create a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include either a maximum of 32 VLAN ranges or a maximum of 32 stacked VLAN ranges, but cannot include a combination of VLAN ranges and stacked VLAN ranges.
- Configuring the NAS-Port-Type attribute and NAS-Port extended format on a per-VLAN, per-stacked VLAN, or per-physical interface basis overrides the global settings for these attributes configured in an access profile.
- If the NAS-Port-Type attribute and the NAS-Port extended format are not configured on a per-VLAN basis (in a 1:1 access model) or on a per-stacked VLAN basis (in an N:1 access model), the router uses

the global settings configured for these attributes in an access profile for all RADIUS request messages.

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

On MX Series routers with MPC/MIC interfaces, you can configure the NAS-Port-Type (61) RADIUS IETF attribute, and an extended format for the NAS-Port (5) RADIUS IETF attribute, on a per-physical interface, per-VLAN, or per-stacked VLAN (S-VLAN) basis. The router passes the NAS-Port-Type and NAS-Port attributes to the RADIUS server during the authentication, authorization, and accounting (AAA) process.

To configure the NAS-Port-Type and NAS-Port extended format on a per-physical interface, per-VLAN, or per-stacked VLAN basis, you must create a NAS-Port options definition, which includes the following components:

- **NAS-Port-Type value**—Specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.
- **NAS-Port extended format**—Configures the number of bits (bit width) for each field in the NAS-Port attribute, which specifies the physical port number of the NAS that is authenticating the subscriber. Fields in the NAS-Port attribute include: slot, adapter, port, VLAN, and S-VLAN. Optionally, you can also use the stacked option as part of the `nas-port-extended-format` statement to include S-VLAN IDs, in addition to VLAN IDs, in the extended format. If you do not configure the stacked option, stacked VLAN IDs are not included in the extended format.
- **VLAN ranges or S-VLAN ranges**—Defines the VLAN range of subscribers or stacked VLAN range of subscribers to which each NAS-Port options definition applies.



NOTE: You can create a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 stacked VLAN ranges, but *cannot* include a combination of VLAN ranges and stacked VLAN ranges.

To configure the NAS-Port-Type and NAS-Port extended format on a per-physical interface, per-VLAN, or per-stacked VLAN basis:

1. Specify the physical interface you want to configure.
2. Enable VLAN tagging, stacked VLAN tagging, or flexible VLAN tagging on the interface.
 - For VLAN tagging, see [Enabling VLAN Tagging](#).
 - For stacked VLAN tagging, see [Configuring Stacked VLAN Tagging](#).
 - For flexible VLAN tagging, also referred to as mixed tagging, see [Enabling VLAN Tagging](#).

3. Specify that you want to configure RADIUS options for a physical interface, VLAN, or S-VLAN.

```
[edit interfaces interface-name]
user@host> edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port-Type, and the VLAN ranges or stacked VLAN ranges to which the named NAS-Port options definition applies.
 - For per-physical interface configurations, see ["Configuring the RADIUS NAS-Port-Type per Physical Interface" on page 702](#).
 - For per-VLAN configurations, see ["Configuring the RADIUS NAS-Port-Type per VLAN" on page 703](#).
 - For per-stacked VLAN configurations, see ["Configuring the RADIUS NAS-Port-Type per Stacked VLAN" on page 705](#).
6. Configure the NAS-Port extended format, and the VLAN ranges or stacked VLAN ranges to which the named NAS-Port options definition applies.
 - For per-physical interface configurations, see ["Configuring the RADIUS NAS-Port Extended Format per Physical Interface" on page 707](#).
 - For per-VLAN configurations, see ["Configuring the RADIUS NAS-Port Extended Format per VLAN" on page 708](#).
 - For per-stacked VLAN configurations, see ["Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN" on page 710](#).

Manual Configuration of the NAS-Port-Type RADIUS Attribute

Subscriber management uses the NAS-Port-Type (RADIUS attribute 61) to identify the type of physical port that is used to authenticate subscribers. By default, subscriber management uses a NAS-Port-Type of ethernet.

You can optionally configure access profiles to provide the value for the NAS-Port-Type attribute, which enables you to explicitly specify the NAS port type that is used for a given connection. For example, you might configure an access profile that specifies that a NAS port type of wireless is used for all Ethernet connections that are managed by that access profile.



NOTE: The **ethernet-port-type-virtual** *configuration statement* takes precedence over the **nas-port-type** statement when you include both statements in the same access profile. When you include the **ethernet-port-type-virtual** statement, subscriber management uses the RADIUS attribute value of 5, which specifies a NAS port type of **virtual**.

Table 9 on page 700 shows the supported port type values for RADIUS attribute 61 (NAS-Port-Type) that you can include in an access profile.

Table 9: RADIUS NAS-Port-Type Values

Statement Option	NAS-Port-Type Value	Description
<i>value</i>	0–65535	Number that indicates either the IANA-assigned value for the RADIUS port type or a custom number-to-port type defined by the user
adsl-cap	12	Asymmetric DSL, carrierless amplitude phase (CAP) modulation
adsl-dmt	13	Asymmetric DSL, discrete multitone (DMT)
async	0	Asynchronous
cable	17	Cable
ethernet	15	Ethernet
fddi	21	Fiber Distributed Data Interface
g3-fax	10	G.3 Fax
hdlc-clear-channel	7	HDLC Clear Channel
iapp	25	Inter-Access Point Protocol (IAPP)

Table 9: RADIUS NAS-Port-Type Values (Continued)

Statement Option	NAS-Port-Type Value	Description
idsl	14	ISDN DSL
isdn-sync	2	ISDN Synchronous
isdn-v110	4	ISDN Async V.110
isdn-v120	3	ISDN Async V.120
piafs	6	Personal Handyphone System (PHS) Internet Access Forum Standard
sdsl	11	Symmetric DSL
sync	1	Synchronous
token-ring	20	Token Ring
virtual	5	Virtual
wireless	18	Other wireless
wireless-1x-ev	24	Wireless 1xEV
wireless-cdma2000	22	Wireless code division multiple access (CDMA) 2000
wireless-ieee80211	19	Wireless 802.11
wireless-umts	23	Wireless universal mobile telecommunications system (UMTS)
x25	8	X.25

Table 9: RADIUS NAS-Port-Type Values (Continued)

Statement Option	NAS-Port-Type Value	Description
x75	9	X.75
xdsl	16	DSL of unknown type

Configuring the RADIUS NAS-Port-Type per Physical Interface

As an alternative to globally configuring the NAS-Port-Type (61) RADIUS attribute in an access profile, you can configure the NAS-Port-Type on a per-physical interface basis as part of a NAS-Port options definition. The NAS-Port-Type specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.

Configuring NAS-Port options definitions on a per-physical interface basis is useful in network configurations that use a 1:1 access model or an N:1 access model.

To configure the NAS-Port-Type RADIUS attribute per physical interface:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a physical interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port-Type.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-type port-type
```

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set vlan-ranges (any | low-tag-high-tag)
```

Per-physical interface configurations typically require you to create a VLAN range that consists of all VLAN IDs on the physical interface. To do so, use the `any` option with the `vlan-ranges` statement.

The following example shows a per-interface NAS-Port options definition named `subscribers-east` that configures the `wireless-umts` NAS-Port-Type for a VLAN range consisting of all VLAN IDs on Gigabit Ethernet physical interface `ge-1/0/0`.

```
[edit interfaces ge-1/0/0 radius-options]
nas-port-options subscribers-east {
  nas-port-type wireless-umts;
  vlan-ranges {
    any;
  }
}
```

Configuring the RADIUS NAS-Port-Type per VLAN

As an alternative to globally configuring the NAS-Port-Type (61) RADIUS attribute in an access profile, you can configure the NAS-Port-Type on a per-VLAN basis as part of a NAS-Port options definition. The NAS-Port-Type specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.

Configuring NAS-Port options definitions on a per-VLAN basis is useful in network configurations that use a 1:1 access model.

To configure the NAS-Port-Type RADIUS attribute per VLAN:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a VLAN interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port-Type.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-type port-type
```

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set vlan-ranges (any | low-tag-high-tag)
```

Per-VLAN configurations typically require you to create a VLAN range that consists of a single VLAN ID on the physical interface. To do so, set the *low-tag* and *high-tag* options in the *vlan-ranges* statement to the same value, as shown in the following example.

The following example shows a per-VLAN NAS-Port options definition named subscribers-west that configures the ethernet NAS-Port-Type for VLAN ID 3 on Gigabit Ethernet physical interface ge-1/1/0.

```
[edit interfaces ge-1/1/0 radius-options]
nas-port-options subscribers-west {
  nas-port-type ethernet;
  vlan-ranges {
    3-3;
  }
}
```

Configuring the RADIUS NAS-Port-Type per Stacked VLAN

As an alternative to globally configuring the NAS-Port-Type (61) RADIUS attribute in an access profile, you can configure the NAS-Port-Type on a per-stacked VLAN basis as part of a NAS-Port options definition. The NAS-Port-Type specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.

Configuring NAS-Port options definitions on a per-stacked VLAN basis is useful in network configurations that use an N:1 access model.

To configure the NAS-Port-Type RADIUS attribute per stacked VLAN:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable stacked VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set stacked-vlan-tagging
```

Setting stacked VLAN tagging enables you to configure dual VLAN tags for all logical interfaces on the physical interface. You must enable stacked VLAN tagging before you can configure the stacked VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a stacked VLAN interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```


4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port-Type.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-type port-type
```

6. Configure the stacked VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any
```

Per-stacked VLAN configurations require you to create a stacked VLAN range of subscribers to which the NAS-Port options definition applies. You must configure the low and high outer tags (VLAN IDs) in the range 1 through 4094, and the inner tag (S-VLAN ID) as any to represent all S-VLAN ID tags.

7. Repeat Steps 3 through 6 to configure additional NAS-Port options definitions on this interface.

The following example creates two NAS-Port options definitions, subscribers-north and subscribers-south, configured on a per-stacked VLAN basis on Gigabit Ethernet physical interface ge-1/1/0.

The subscribers-north definition configures a NAS-Port-Type user-defined value (4711) for a stacked VLAN range with outer VLAN ID 1 and all inner S-VLAN IDs. The subscribers-south definition configures a NAS-Port-Type user-defined value (4722) for a stacked VLAN range with outer VLAN IDs in the range 2 through 10, and all inner S-VLAN IDs.

```
[edit interfaces ge-2/0/1 radius-options]
nas-port-options subscribers-north {
  nas-port-type 4711;
  stacked-vlan-ranges {
    1-1,any;
  }
}
nas-port-options subscribers-south {
  nas-port-type 4722;
  stacked-vlan-ranges {
    2-10,any;
```

```
}
}
```

Configuring the RADIUS NAS-Port Extended Format per Physical Interface

As an alternative to globally configuring the extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-physical interface basis as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN.

Configuring NAS-Port options definitions on a per-physical interface basis is useful in network configurations that use a 1:1 access model or an N:1 access model.

To configure an extended format for the NAS-Port RADIUS attribute per physical interface:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a physical interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port extended format.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-extended-format slot-width width adapter-width width port-width width
vlan-width width
```

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set vlan-ranges (any | low-tag-high-tag)
```

Per-physical interface configurations typically require you to create a VLAN range that consists of all VLAN IDs on the physical interface. To do so, use the `any` option with the `vlan-ranges` statement.

The following example shows a per-interface NAS-Port options definition named `boston-subscribers` that configures a NAS-Port extended format consisting of an 8-bit slot field, 8-bit adapter field, 8-bit port field, and 4-bit VLAN field. The `boston-subscribers` definition applies to a VLAN range consisting of all VLAN IDs on Gigabit Ethernet physical interface `ge-2/0/1`.

```
[edit interfaces ge-2/0/1 radius-options]
nas-port-options boston-subscribers {
  nas-port-extended-format {
    slot-width 8;
    adapter-width 8;
    port-width 8;
    vlan-width 4;
  }
  vlan-ranges {
    any;
  }
}
```

Configuring the RADIUS NAS-Port Extended Format per VLAN

As an alternative to globally configuring the extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-VLAN basis as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN.

Configuring NAS-Port options definitions on a per-VLAN basis is useful in network configurations that use a 1:1 access model.

To configure an extended format for the NAS-Port RADIUS attribute per VLAN:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a VLAN interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port extended format.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-extended-format slot-width width adapter-width width port-width width
vlan-width width
```

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set vlan-ranges (any | low-tag-high-tag)
```

Per-VLAN configurations typically require you to create a VLAN range that consists of a single VLAN ID on the physical interface. To do so, set the *low-tag* and *high-tag* options in the *vlan-ranges* statement to the same value, as shown in the following example.

The following example shows a per-VLAN NAS-Port options definition named `paris-subscribers` that configures a NAS-Port extended format consisting of a 4-bit slot field, 2-bit adapter field, 4-bit port field, and 2-bit VLAN field. The `paris-subscribers` definition applies to VLAN ID 1 on Gigabit Ethernet physical interface `ge-1/0/1`.

```
[edit interfaces ge-1/0/1 radius-options]
nas-port-options paris-subscribers {
  nas-port-extended-format {
    slot-width 4;
    adapter-width 2;
    port-width 4;
    vlan-width 2;
  }
  vlan-ranges {
    1-1;
  }
}
```

Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN

As an alternative to globally configuring the extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per- stacked VLAN basis as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN.

Configuring NAS-Port options definitions on a per-stacked VLAN basis is useful in network configurations that use an N:1 access model.

To configure an extended format for the NAS-Port RADIUS attribute per stacked VLAN:

1. Specify the interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Enable stacked VLAN tagging on the interface.

```
[edit interfaces interface-name]
user@host# set stacked-vlan-tagging
```

Setting stacked VLAN tagging enables you to configure dual VLAN tags for all logical interfaces on the physical interface. You must enable stacked VLAN tagging before you can configure the stacked VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a stacked VLAN interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

4. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

5. Configure the NAS-Port extended format.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-extended-format slot-width width adapter-width width port-width width
vlan-width width stacked
```

To include S-VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format, include the stacked option in the nas-port-extended-format statement.

6. Configure the stacked VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any
```

Per-stacked VLAN configurations require you to create a stacked VLAN range of subscribers to which the NAS-Port options definition applies. You must configure the low and high outer tags (VLAN IDs) in the range 1 through 4094, and the inner tag (S-VLAN ID) as any to represent all S-VLAN ID tags.

7. Repeat Steps 3 through 6 to configure additional NAS-Port options definitions on this interface.

The following example creates two NAS-Port options definitions, *chicago-subscribers* and *barcelona-subscribers*, configured on a per-stacked VLAN basis on Gigabit Ethernet physical interface ge-3/2/1.

The *chicago-subscribers* definition configures a NAS-Port extended format consisting of a 8-bit slot field, 8-bit adapter field, 8-bit port field, 4-bit stacked VLAN field, and 4-bit VLAN field. Because the stacked option is configured in this definition, S-VLAN IDs, in addition to VLAN IDs, are included in the extended format. The *chicago-subscribers* definition applies to a stacked VLAN range with outer VLAN ID 1, and all inner S-VLAN IDs.

The `barcelona-subscribers` definition configures a NAS-Port extended format consisting of a 8-bit slot field, 8-bit adapter field, 8-bit port field, 4-bit stacked VLAN field, and 4-bit VLAN field. Because the stacked option is *not* configured in this definition, S-VLAN IDs are not included in the extended format. The `barcelona-subscribers` definition applies to a stacked VLAN range with outer VLAN IDs in the range 2 through 10, and all inner S-VLAN IDs.

```
[edit interfaces ge-3/2/1 radius-options]
nas-port-options chicago-subscribers {
  nas-port-extended-format {
    slot-width 8;
    adapter-width 8;
    port-width 8;
    stacked-vlan-width 4;
    vlan-width 4;
    stacked;
  }
  stacked-vlan-ranges {
    1-1,any;
  }
}

nas-port-options barcelona-subscribers {
  nas-port-extended-format {
    slot-width 8;
    adapter-width 8;
    port-width 8;
    stacked-vlan-width 4;
    vlan-width 4;
  }
  stacked-vlan-ranges {
    2-10,any;
  }
}
```

Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces

As an alternative to globally configuring an extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-physical interface basis for both Ethernet subscribers and ATM subscribers as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field of the NAS-Port attribute, including: slot, adapter, port, ATM virtual path identifier (VPI), and ATM virtual circuit identifier (VCI).

To configure the NAS-Port extended format for an ATM interface, include one or both of the following options in the `nas-port-extended-format` statement along with the other options as appropriate for your needs:

- `vpi-width`—Number of bits in the ATM VPI field, in the range 1 through 32
- `vci-width`—Number of bits in the ATM VCI field, in the range 1 through 32



NOTE: For ATM subscribers, the combined total of the widths of all fields must not exceed 32 bits, or the configuration fails. The router may truncate the values of individual fields depending on the bit width you specify.

To configure an extended format for the NAS-Port RADIUS attribute for an ATM interface:

1. Specify the ATM interface you want to configure.

```
[edit]
user@host# edit interfaces interface-name
```

2. Specify that you want to configure RADIUS options for a physical interface.

```
[edit interfaces interface-name]
user@host# edit radius-options
```

3. Create a named NAS-Port options definition.

```
[edit interfaces interface-name radius-options]
user@host# edit nas-port-options nas-port-options-name
```

4. Configure the NAS-Port extended format.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set nas-port-extended-format slot-width width adapter-width width port-width width
vpi-width width vci-width width
```

The following example shows a NAS-Port options definition named `boston-subscribers` for ATM interface `at-1/0/4` that configures a NAS-Port extended format with an ATM slot width of 6 bits, ATM

adapter width of 3 bits, ATM port width of 4 bits, ATM VPI width of 12 bits, and ATM VCI width of 24 bits.

```
[edit interfaces at-1/0/4 radius-options]
nas-port-options boston-subscribers {
  nas-port-extended-format {
    slot-width 6;
    adapter-width 3;
    port-width 4;
    vpi-width 12;
    vci-width 24;
  }
}
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, a logical port number is added to the default format for only channelized interfaces.

RELATED DOCUMENTATION

- [Configuring Access Profile Options for Interactions with RADIUS Servers | 653](#)
- [RADIUS Servers and Parameters for Subscriber Access | 646](#)

RADIUS Logical Line Identification

IN THIS SECTION

- [RADIUS Logical Line Identifier \(LLID\) Overview | 715](#)
- [RADIUS Attributes for LLID Preauthentication Requests | 716](#)
- [Configuring Logical Line Identification \(LLID\) Preauthentication | 718](#)

- [Configuring a Port and Password for LLID Preauthentication Requests | 719](#)
- [Verifying and Managing LLID Preauthentication Configuration | 721](#)

RADIUS Logical Line Identifier (LLID) Overview

The logical line identification (LLID) feature helps service providers maintain a reliable and up-to-date customer database for those subscribers who frequently move from one physical line to another. The LLID is designed to provide the service provider with a configurable calling station ID for the subscriber access line. A calling station ID is derived from the physical line location and the subscriber client's information. The line information derived from the facility of the service provider is not friendly for the access line wholesaler to manage access line ownership when subscribers frequently move physical locations. The LLID feature is based on a virtual port — the LLID — rather than the physical line used by the subscriber. The LLID provides AAA driven line information management with a service provider (usually a wholesaler).

The LLID is an alphanumeric string that is based on the subscriber user name and circuit ID. The LLID logically identifies the subscriber line, and is mapped to the subscriber's physical line in the service provider customer database. When the subscriber moves to a different location and different physical line, the database is updated to map the LLID to the new physical line. Because the subscriber's LLID remains constant, it provides service providers with a secure and reliable means for tracking subscribers and maintaining an accurate customer database. Subscriber management supports the LLID feature for PPP subscribers over PPPoE, PPPoA, and LAC.

To assign an LLID to a subscriber, the router issues two RADIUS access requests. The first request is a preauthentication request, which obtains the LLID from a RADIUS preauthentication server. The second request is the standard authentication request sent to the RADIUS authentication server.

The following sequence of steps describes how subscriber management obtains and uses the LLID. The procedure assumes that preauthentication is enabled on the router and that the RADIUS preauthentication and authentication servers are configured.

1. The PPP subscriber sends an Authentication-Request message to the router.
2. The router sends an Access-Request message to the RADIUS preauthentication server to obtain an LLID for the subscriber.
3. The preauthentication server returns the LLID to the router in the Calling-Station-Id attribute (RADIUS attribute 31) in the Access-Accept message.



NOTE: This step includes a non-standard use of the Calling-Station-Id attribute. This attribute is typically present in RADIUS request messages, such as an Access-Request, not in response messages. Also, the router ignores all RADIUS attributes, other than the Calling-Station-Id, that are returned in the preauthentication Access-Accept message. In addition, any **radius options** that are configured on the router, such as **calling-station-id-format**, have no effect on the Calling-Station-Id attribute in the preauthentication request.

- 4. The router encodes the Calling-Station-Id (the LLID) in a second Access-Request message and sends the message to the RADIUS authentication server. This authentication request is the standard use of the Calling-Station-Id attribute.
- 5. The RADIUS authentication server returns an Access-Accept message to the router. The Access-Accept message includes attributes for the subscriber session.



NOTE: Once the preauthenticated subscriber has been successfully authenticated by the RADIUS authentication server, all subsequent RADIUS request messages, such as Accounting-Request messages, will include the LLID in the Calling-Station-Id attribute.



NOTE: For tunneled PPP subscribers, the router, acting as an L2TP access concentrator (LAC), encodes the LLID into Calling Number AVP (L2TP attribute 22) and sends the attribute to the L2TP network server (LNS) in an Incoming-Call-Request (ICRQ) packet. After a successful preauthentication request, the router always encodes the LLID in the L2TP Calling Number AVP.

RADIUS Attributes for LLID Preauthentication Requests

Table 10 on page 717 lists the RADIUS IETF attributes used in a preauthentication request to obtain a subscriber's LLID, and describes the information that is included in the attributes. In some cases, preauthentication uses an attribute for information that is different than the IETF description—the table indicates any non-standard use of RADIUS attributes.

Table 10: RADIUS Attributes for LLID Preauthentication Requests

Attribute Number	Attribute Name	Description
1	User-Name	<p>(Non-standard use of attribute.) Identifying information for the user associated with the LLID, in the following format.</p> <p><i>nas-port: nas-ip-address: nas-port-id</i></p> <p>Example: nas-port:198.51.100.117:ge-1/0/5:100</p> <p>NOTE: The router strips any dynamically generated information from the User-Name attribute during preauthentication.</p>
2	User-Password	<p>(Non-standard use of attribute.) Password of the user to be authenticated.</p> <p>Example: Always set to juniper</p>
4	NAS-IP-Address	<p>IP address of the network access server (NAS) that is requesting authentication of the user</p> <p>Example: 198.51.100.117</p>
5	NAS-Port	Physical port number of the NAS that is authenticating the user. Always interpreted as a bit field
6	Service-Type	<p>Type of service the user requested or the type of service to be provided.</p> <p>Example: gold-service</p>
61	NAS-Port-Type	Type of physical port the NAS is using to authenticate the user. You can use the ethernet-port-type-virtual statement to configure this to virtual (type 5).
77	Connect-Info	<p>(Non-standard use of attribute.) The user name.</p> <p>Example: jdoe@xyzcorp.example.com</p>

Table 10: RADIUS Attributes for LLID Preauthentication Requests (*Continued*)

Attribute Number	Attribute Name	Description
87	NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user. Includes any dynamically generated information. Example: ge 1/0/5:100

Configuring Logical Line Identification (LLID) Preauthentication

The logical line identification (LLID) feature enables service providers to track subscribers on the basis of a virtual port — the LLID — rather than by the physical port used by the subscriber. The LLID is assigned by a RADIUS preauthentication server, which you configure in an access profile.

To configure the router to support preauthentication for the LLID feature:



NOTE: You cannot configure the preauthentication statements in this procedure if you have configured the `radius attributes exclude` statement to exclude the Calling-Station-ID attribute from RADIUS Access-Request messages.

1. Specify the access profile you want to use for the subscriber preauthentication support.

```
[edit]
user@host# edit access profile profile-name
```

2. Specify the order in which the router uses the supported preauthentication methods. **radius** is the only supported authentication method.

```
[edit access profile profile-name]
user@host# set preauthentication-order radius
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile profile-name]
user@host# edit radius
```

- Specify the IP address of the RADIUS server used for preauthentication.

```
[edit access profile profile-name radius]
user@host# set preauthentication-server 192.168.100.10
```



NOTE: The preauthentication feature uses the `retry` and `timeout` parameters that are configured for the RADIUS authentication server.

- (Optional) Display AAA preauthentication statistics.

```
user@host> show network-access aaa statistics preauthentication
Preauthentication module statistics
Requests received: 2118
Multistack requests: 0
Accepts: 261
Rejects: 975
Challenges: 0
Requests timed out: 882
```

- (Optional) Verify configuration of the RADIUS preauthentication server.

```
user@host1> show radius pre-authentication servers
```

```

RADIUS Pre-Authentication Configuration
-----
      Udp  Retry      Maximum  Dead
IP Address  Port  Count  Timeout Sessions  Time  Secret
-----
203.0.113.168 1812  3      3      255      0    radius
```

Configuring a Port and Password for LLID Preauthentication Requests

You can configure a router that operates as the RADIUS client to contact a RADIUS server for authentication and preauthentication requests on two different UDP ports and using different secret passwords. Similar to configuring the port numbers for authentication and accounting requests, you can define a unique port number that the router uses to contact the RADIUS server for logical line

identification (LLID) preauthentication requests. You can also define a unique password for preauthentication requests. If you do not configure a separate UDP port or secret for preauthentication purposes, the same UDP port and secret that you configure for authentication messages is used.

To configure a unique UDP port number to be used to contact the RADIUS server for preauthentication requests, include the `preauthentication-port port-number` statement at the `[edit access radius-server server-address]` or `[edit access profile profile-name radius-server server-address]` hierarchy level.

- To specify the UDP port for all of the access profiles:

```
[edit access]
radius-server server-address {
    preauthentication-port port-number;
}
```

- To specify the UDP port for a specific access profile:

```
[edit access]
profile profile-name {
    radius-server server-address {
        preauthentication-port port-number;
    }
}
```

To configure the password to be used to contact the RADIUS preauthentication server, include the `preauthentication-secret password` statement at the `[edit access radius-server server-address]` or `[edit access profile profile-name radius-server server-address]` hierarchy level.

- To specify the password for all of the access profiles:

```
[edit access]
radius-server server-address {
    preauthentication-secret password;
}
```

- To specify the password for a specific access profile:

```
[edit access]
profile profile-name {
    radius-server server-address {
```

```

    preauthentication-secret password;
  }
}

```

Verifying and Managing LLID Preauthentication Configuration

IN THIS SECTION

- [Purpose | 721](#)
- [Action | 721](#)

Purpose

Display statistics and configuration information related to logical line identification (LLID) preauthentication.

Action

- To display LLID preauthentication statistics:

```
user@host> show network-access aaa statistics preauthentication
```

- To display information about preauthentication servers:

```
user@host> show network-access aaa radius-servers
```

RADIUS Authentication and Accounting Basic Configuration

IN THIS SECTION

- [Configuring Authentication and Accounting Parameters for Subscriber Access | 722](#)
- [Specifying the Authentication and Accounting Methods for Subscriber Access | 722](#)

- [Specifying RADIUS Authentication and Accounting Servers for Subscriber Access | 723](#)
- [Configuring Local Authentication and Authorization for Subscribers | 724](#)

Configuring Authentication and Accounting Parameters for Subscriber Access

You use an access profile to configure authentication and accounting support for the subscriber access management feature. The access profile enables you to specify the type of methods used for authentication and accounting. You can also configure how subscriber access management collects and uses accounting statistics.

To configure authentication and accounting for subscriber access:

1. Specify the authentication and accounting methods to use.
See ["Specifying the Authentication and Accounting Methods for Subscriber Access" on page 722](#).
2. Specify how accounting statistics are collected.
See ["Configuring Per-Subscriber Session Accounting" on page 744](#).

Specifying the Authentication and Accounting Methods for Subscriber Access

You can specify the authentication and accounting methods that subscriber access management uses.

You can configure multiple authentication and accounting methods—the `authentication-order` and `accounting order` statements specify the order in which the subscriber access management feature uses the methods. For example, an authentication entry of `radius password` specifies that RADIUS authentication is performed first; if it times out (for example, if the RADIUS server is unreachable), then local authentication (`password`) is attempted. However, if a method rejects the authentication attempt, no subsequent method is attempted. If `password` is configured as the first method to be attempted, authentication is always either accepted or rejected; in either case, no other method is attempted.

You can specify the following authentication methods with the `authentication-order` statement:

- `radius`—RADIUS-based authentication using an external RADIUS server.
- `password`—Local authentication using locally configured and stored usernames and passwords.

Subscriber access management does not support the `password` option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, you can use the `password` option to provide local authentication for individual subscribers, typically when you do not have external authentication and authorization servers, or when you want to use local authentication as a backup to external authentication. In this case, you configure the actual subscriber password with the `password` option of the `subscriber username` statement in the access profile. In earlier releases you must always specify the `radius` authentication method.

You can specify the following accounting methods:

- radius—RADIUS-based accounting using an external RADIUS server.

To configure the authentication and accounting methods for subscriber access management:

1. Specify the authentication methods and the order in which they are used.

```
[edit access profile profile-name]
user@host# set authentication-order method
```

2. Specify the accounting method.

```
[edit access profile profile-name]
user@host# set accounting order radius
```

Specifying RADIUS Authentication and Accounting Servers for Subscriber Access

You can specify one or more RADIUS authentication or accounting servers to use for subscriber access management.

To configure RADIUS authentication and accounting support:

1. Specify that you want to configure RADIUS support.

```
[edit access profile isp-bos-metro-fiber-basic]
user@host# edit radius
```

2. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile isp-bos-metro-fiber-basic radius]
user@host# set authentication-server 192.168.1.251
```

3. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile isp-bos-metro-fiber-basic radius]
user@host# set accounting-server 192.168.1.250
```

To configure multiple RADIUS authentication or accounting servers:

- Specify the IP addresses of all RADIUS servers used for authentication or accounting.

```
[edit access profile isp-bos-metro-fiber-basic radius]
user@host# set authentication-server 192.168.1.251 192.168.1.252
user@host# set accounting-server 192.168.1.250 192.168.1.251
```

Configuring Local Authentication and Authorization for Subscribers

Starting in Junos OS Release 18.2R1, you can configure local authentication and limited local authorization for subscribers. Local authentication supports all subscriber types that are currently supported by subscriber management and services on MX Series routers. Local authentication and authorization is useful in the following circumstances:

- When you do not want to use external authentication and authorization servers.
- When you want local authentication and authorization to provide a backup method in the event RADIUS authentication fails.
- When you are migrating a network from E Series routers running JunosE software to MX Series routers running Junos OS.

Enable local authentication and authorization for subscribers by configuring the `password` option to be configured as an `authentication-order` method for the access profile. Then configure a password for each subscriber you want to authenticate locally. When a subscriber associated with the access profile logs in, the login username is compared to the configured username. If that matches, then the login password is compared to the configured password. Local authentication failures result from credential mismatches; that is, either the subscriber username or password do not match.

Local authentication can take the form of either of the following:

- User password authentication—The configured password is used to verify the subscriber's login password.
- Challenge handshake authentication (CHAP)—The configured password acts as the challenge secret to verify the subscriber's challenge password and challenge response credential.

You can also optionally configure several attributes, such as address pool, logical system, or routing instance, to be authorized locally for the subscriber when authentication is successful. If you do not configure an address or address pool for local authorization, address assignment is based on network matching or the first address pool assigned to the routing instance.



NOTE: Local authentication and authorization support a chassis-wide maximum of 100 subscribers. If subscribers are configured in access profiles where authentication-order password is not configured, local authentication does not occur, but these subscribers count against the system limit of 100 subscribers for local authentication.

To configure local authentication and authorization:

1. Enable local authentication.

```
[edit access profile profile-name]
user@host# set authentication-order password
```

If you want only local authentication to be used, then configure `password` as the only authentication method. If you want local authentication to back up RADIUS authentication in the event the method times out, then you must configure `radius` as the first method and `password` as the second method, like so:

```
[edit access profile profile-name]
user@host# set authentication-order [radius password]
```

If you configure `password` as the first method, authentication is always either accepted or rejected. In either case, a second method is never attempted.

2. Configure the local password for the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username password password
```

3. (Optional) Configure an IPv4 address for the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username framed-ip-address ipv4-address
```

4. (Optional) Configure an address pool to assign an IPv4 address for the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username framed-pool ipv4-pool-name
```

5. (Optional) Configure an address pool to assign a router advertisement IPv6 prefix or a DHCPv6 IA_NA/128 address for the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username framed-ipv6-pool ipv6-pool-name
```

6. (Optional) Configure an address pool to locally allocate a delegated IPv6 prefix.

```
[edit access profile profile-name]
user@host# set subscriber username delegated-pool delegated-pool-name
```

7. (Optional) Configure a logical system and if desired a routing instance assigned to the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username target-logical-system logical-system-name <target-routing-
instance (default | routing-instance-name)>
```

8. (Optional) Configure a routing instance for the subscriber.

```
[edit access profile profile-name]
user@host# set subscriber username target-routing-instance (default | routing-instance-name)
```

You can use the following `show` commands to display information about local authentication:

- `show network-access aaa statistics authentication detail`—Displays failure statistics for local authentication.
- `show network-access requests statistics`—Displays both local authentication and local reauthentication statistics such as requests received and the number of success and failure responses.
- `show network-access aaa statistics re-authentication`—Displays reauthentication statistics, but they are aggregated from both local authentication and RADIUS.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.2R1	Starting in Junos OS Release 18.2R1, you can configure local authentication and limited local authorization for subscribers.

RELATED DOCUMENTATION

[AAA Service Framework Overview | 2](#)

[RADIUS Servers and Parameters for Subscriber Access | 646](#)

[Configuring Local Authentication and Authorization for Subscribers | 724](#)

[Configuring Per-Subscriber Session Accounting | 744](#)

[Specifying the Authentication and Accounting Methods for Subscriber Access | 722](#)

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers

IN THIS SECTION

- [Benefits of Reauthentication | 728](#)
- [Functionality | 728](#)
- [Dual-Stack Subscribers | 731](#)
- [Packet Flow | 733](#)
- [RADIUS Attributes Supported for Reauthentication | 736](#)

RADIUS Change of Authorization (CoA) messages, specified in RFC 5176, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*, are used to activate or deactivate client services and to change certain client session characteristics without logging out the client, thus avoiding interruption to the subscriber. In some circumstances, it may be preferable to use *reauthentication* of the subscriber as the method to alter client session services and characteristics without interruption.

For example, the following customer deployment modes both require changes in attributes during the life of a session.

- Residential subscribers—Residential subscribers may change service plans throughout the life of a session by online service selection or direct calling to the provider. The change in service plan is propagated to the DHCP local server by changing the value of the DHCP client Agent Remote ID. The Agent Remote ID is conveyed in option 82, suboption 2, for DHCPv4 clients and in option 37 for DHCPv6 clients.

When reauthentication is configured, the change in service plan is detected, triggering reauthentication; the new service plan and any changed attributes are returned by the RADIUS server and implemented for the subscriber.

- **Business subscribers**—Business subscribers often need changes in attributes (particularly framed routes) during any given session. The desired change in attributes is not initiated by a change in service plans.

When reauthentication is configured, negotiation of the lease renewal triggers reauthentication. Any changes in attributes or services are provided in the Access-Accept message from the RADIUS server and implemented for the subscriber.

Two alternatives to using reauthentication can both change many more session characteristics than are possible with reauthentication. CoA requests change characteristics without disrupting the subscriber. Logging the subscriber out and then back in can change many more session characteristics but is obviously disruptive.

Benefits of Reauthentication

- Update or modify subscriber session attributes and service plans without using a CoA request.
- Simplify activation of services resulting from frequent subscriber-initiated changes.
- Enable reauthentication per family in dual-stack, single-session configurations.
- Control reauthentication through CLI configuration or a RADIUS VSA.

Functionality

Reauthentication is supported for both DHCPv4 and DHCPv6. It can be triggered when the DHCP local server receives a renew, rebind, discover, or solicit message from a DHCP client. The discover and solicit messages support reauthentication starting in Junos OS Release 18.1R1. Support for the discover and solicit messages means that if a CPE with a bound client reboots and the client sends one of those messages to bring the session back up, reauthentication enables authd to obtain any updates that have been made for the subscriber.

Reauthentication behavior is determined as follows:

- The `reauthenticate lease-renewal` statement specifies reauthentication is triggered when any of the four supported messages is received.
- The `reauthenticate remote-id-mismatch` statement specifies reauthentication is triggered only when the received message includes a change in the value of the DHCP client's Agent Remote ID. The attribute value includes the name of the subscriber service plan, so a change in value signifies a change in service for the subscriber.

- The Juniper Networks `reauthentication-on-renew` VSA (26-206), when returned with a value of 1 in the Access-Accept message from the RADIUS server for the subscriber at login, triggers reauthentication on receipt of any of the four messages. A value of 0 disables reauthentication. The VSA value is stored in the session database whenever it is received. After this VSA has enabled reauthentication, it is checked at each reauthentication attempt. If the value has changed to 0—that is, if a subsequent Access-Accept returned the VSA with a value of 0—the reauthentication process stops for that subscriber.

The CLI configuration (`reauthenticate` statement) and Reauthenticate-On-Renew behavior is additive. Disabling reauthentication with the VSA has an effect only when the `reauthenticate` statement is not configured. When the `reauthenticate` statement is configured with either option, it overrides a VSA value of 0. In the absence of the CLI configuration, the VSA can enable reauthentication by itself.

The reauthentication process is almost identical to the original authentication process. When reauthentication is triggered, the `jdhcpd` process on the local server submits an authentication request to `authd`, which in turn submits an Access-Request message to the RADIUS server to request a second authentication.



NOTE: The reauthentication request fails for any authentication order other than `radius` or `none`. The `authd` process returns a negative acknowledgment (NAK) for any such request.

This Access-Request includes RADIUS state and class attributes that were returned in the original Access-Accept message. These attributes enable the RADIUS server to distinguish the reauthentication request from login (authentication) requests.

The RADIUS server returns an Access-Accept message to `authd` with new attributes for the subscriber. The `authd` process sends an acknowledgment (ACK) with the changes to `jdhcpd`, which sends a DHCP offer to the DHCP client with changed attributes. The DHCP negotiation continues as usual, as shown in [Figure 1 on page 734](#), and the subscriber session continues with the new attribute values. When the reauthentication includes a change in service plan, the RADIUS server returns the new plan with any other changed attributes, if it accepts the request, as shown in [Figure 2 on page 735](#). If the CPE hosting the DHCP client reboots during the process of changing a service plan, reauthentication with the new plan is supported with no disruption in service.

If the RADIUS server rejects a reauthentication request or times out, `authd` sends a NAK to `jdhcpd`, which reviews the included error code. If the error code indicates a timeout, `jdhcpd` sends an ACK to the DHCP client and the subscriber session is maintained with the original attributes and service. For any other error code, `jdhcpd` sends a DHCPv4 NAK or DHCPv6 REPLY (with the lifetime value set to 0) as a logical NAK, initiates subscriber logout, and deletes the subscriber from the session database.

[Table 11 on page 730](#) describes how `authd` processes requests when a different request type is already in progress for the same subscriber.

Table 11: Processing Multiple Request Types

Request in Progress	Additional Request Received for Same Subscriber	Action
Reauthentication	CoA	authd responds to the CoA with a NAK.
CoA	Reauthentication	authd queues the reauthentication request until the CoA is processed, then processes the reauthentication request.
Reauthentication	Disconnect	authd responds to the disconnect with a NAK.
Disconnect	Reauthentication	authd responds to the reauthentication request with a NAK and continues logging out the subscriber.



BEST PRACTICE: Because the network family does not terminate or reinitiate as part of reauthentication, the subscriber content is not reevaluated with regards to subscriber secure policy mirroring. Do not use as a trigger for subscriber secure policy mirroring any attribute that can change during reauthentication processing.

When reauthentication results in a change in a DHCPv6 subscriber's IP or IPv6 address after a client is bound, the DHCPv6 server evaluates the address change request. The server returns a status code to the client in the identity association (IA) of the reply PDU. Starting in Junos OS Release 18.4R1, when the DHCPv6 server discovers an issue with the address, a status code for NotOnLink is supported in addition to the previously supported codes for NoAddrsAvail and NoPrefixAvail. These status codes are defined as follows:

- NoAddrsAvail (2)—The server cannot assign any addresses for the IA in the client request. It returns the IA with no addresses and NoAddrsAvail.
- NotOnLink (4)—The server determines that the prefix for one or more addresses in any IA in the client request is not appropriate for the link connecting to the client. This code is also used in the event of a reauthentication failure (RADIUS Access-Reject).
- NoPrefixAvail (6)—The server has no prefixes available for the IA in the client request.

If the client receives the NotOnLink status code, it can send another request without any addresses or it can restart the negotiation process. If it sends a request, the DHCPv6 local server ignores the request, expecting a new renegotiation to begin.

Dual-Stack Subscribers

In releases earlier than Junos OS Release 18.1R1, dual-stack DHCP subscribers are treated as independent client sessions. Each stack renews and obtains new services independently.

Starting in Junos OS Release 18.1R1, per-family authentication and reauthentication are supported for dual-stack, single-session subscribers. A dual-stack, single-session subscriber is typically a household with its own VLAN in a 1:1 access model. The household is represented as a single subscriber with a single session in the session database, but it has two separate DHCP bindings, one for each family, DHCPv4 and DHCPv6. Consequently, authd sends separate Access-Requests as each family in the session logs in or attempts to reauthenticate:

- Per-family authentication occurs when a discover or solicit message is received while a subscriber session is in the DHCP init state.
- Per-family reauthentication occurs when reauthentication and on-demand address allocation are both configured and a renew, rebind, discover, or solicit message is received for the family session while it is in the DHCP bound state.



NOTE: On-demand address allocation causes an address to be allocated separately for each family as it logs in. On-demand address allocation must be configured for dual-stack, single-session subscribers or per-family authentication and reauthentication cannot be enabled. For reauthentication, this is true whether it is configured in the CLI or with the Reauthenticate-On-Renew VSA (26-206).

Authentication and reauthentication are both processed per family. The first family to trigger the process is attended to before the other (second) family triggers authentication or reauthentication. Messages from the second family are ignored until the first family is bound. Then the second family request is processed.

If only one family of the dual-stack single session logs in, then only one authentication is processed. Reauthentications are processed for only the one client family.

The authd process classifies attributes as belonging to the DHCPv4 or DHCPv6 family and tags them accordingly. For both authentication and reauthentication, depending on which family initiates the request, authd includes either the DHCP-Options VSA (26-55) or the DHCPv6-Options VSA (26-65). Depending on its configuration, the RADIUS server might return information for only the family that initiated the request (the *requesting family*) or for both families.

When authd receives attributes in the Access-Accept message, the family tags enable authd to determine which attributes correspond to the requesting family or the other (*nonrequesting*) family. Only attributes for the requesting family are written to the session database.

For reauthentication requests, authd compares the returned attributes to the session database to determine whether any changes were made on the RADIUS server. Again, only changes that correspond to the requesting family are written to the session database, overwriting the old values.

Changes during reauthentication are processed as follows:

- **Attribute (other than address)**—When authd determines that one or more of these attributes has changed for the requesting family, it stores the new values in the session database. After authd notifies jdhcpd, it sends an ACK to the DHCP client with the new attribute values.
- **Address or address pool**—When authd detects a change for the requesting family, it notifies the DHCP local server, which in turn sends a NAK (DHCPv4) or logical NAK (DHCPv6) to the DHCP client.

If the requesting family is the only family that is bound, jdhcpd gracefully logs out the subscriber. If the nonrequesting family is also bound, jdhcpd deactivates the requesting family, but leaves the nonrequesting family binding intact, with no disruption in service to the nonrequesting family. The deactivation of the requesting family has no effect on a subsequent triggering of reauthentication by the nonrequesting family.

When the deactivated family subsequently sends a discover or solicit message to log back in, an Access-Request is sent for reauthentication as usual, and the new address received in the Access-Accept is applied to the subscriber.

If the RADIUS server responds to an authentication or reauthentication request with an Access-Reject message, authd notifies the DHCP local server, which in turn sends a NAK (DHCPv4) or logical NAK (DHCPv6) to the DHCP client. The requesting family is terminated gracefully; the family is deactivated and the subscriber is logged out. Then the nonrequesting family is deactivated and logged out, but the client is not notified about the termination.

If liveness detection is running on the nonrequesting family, the client detects loss of connection when the family is terminated, and subsequently sends a discover or solicit message to the DHCP local server. However, if liveness detection is not running, the client does not detect loss of connection until the rebinding time (T2, option 59) expires and service is lost. Depending on the duration of the lease, this could take a long time.



BEST PRACTICE: Configure liveness detection for both address families to reduce the time to detect loss of connection. See [DHCP Liveness Detection Overview](#) for information about configuring liveness detection.

Packet Flow

The following figure describes the sequence for initial negotiation of a subscriber session between a DHCP client, a DHCP server, and the RADIUS server. The client's service plan is specified in the second substring in the remote ID contained in DHCPv4 option 82, suboption 2, or DHCPv6 option 37.

Initial Negotiation

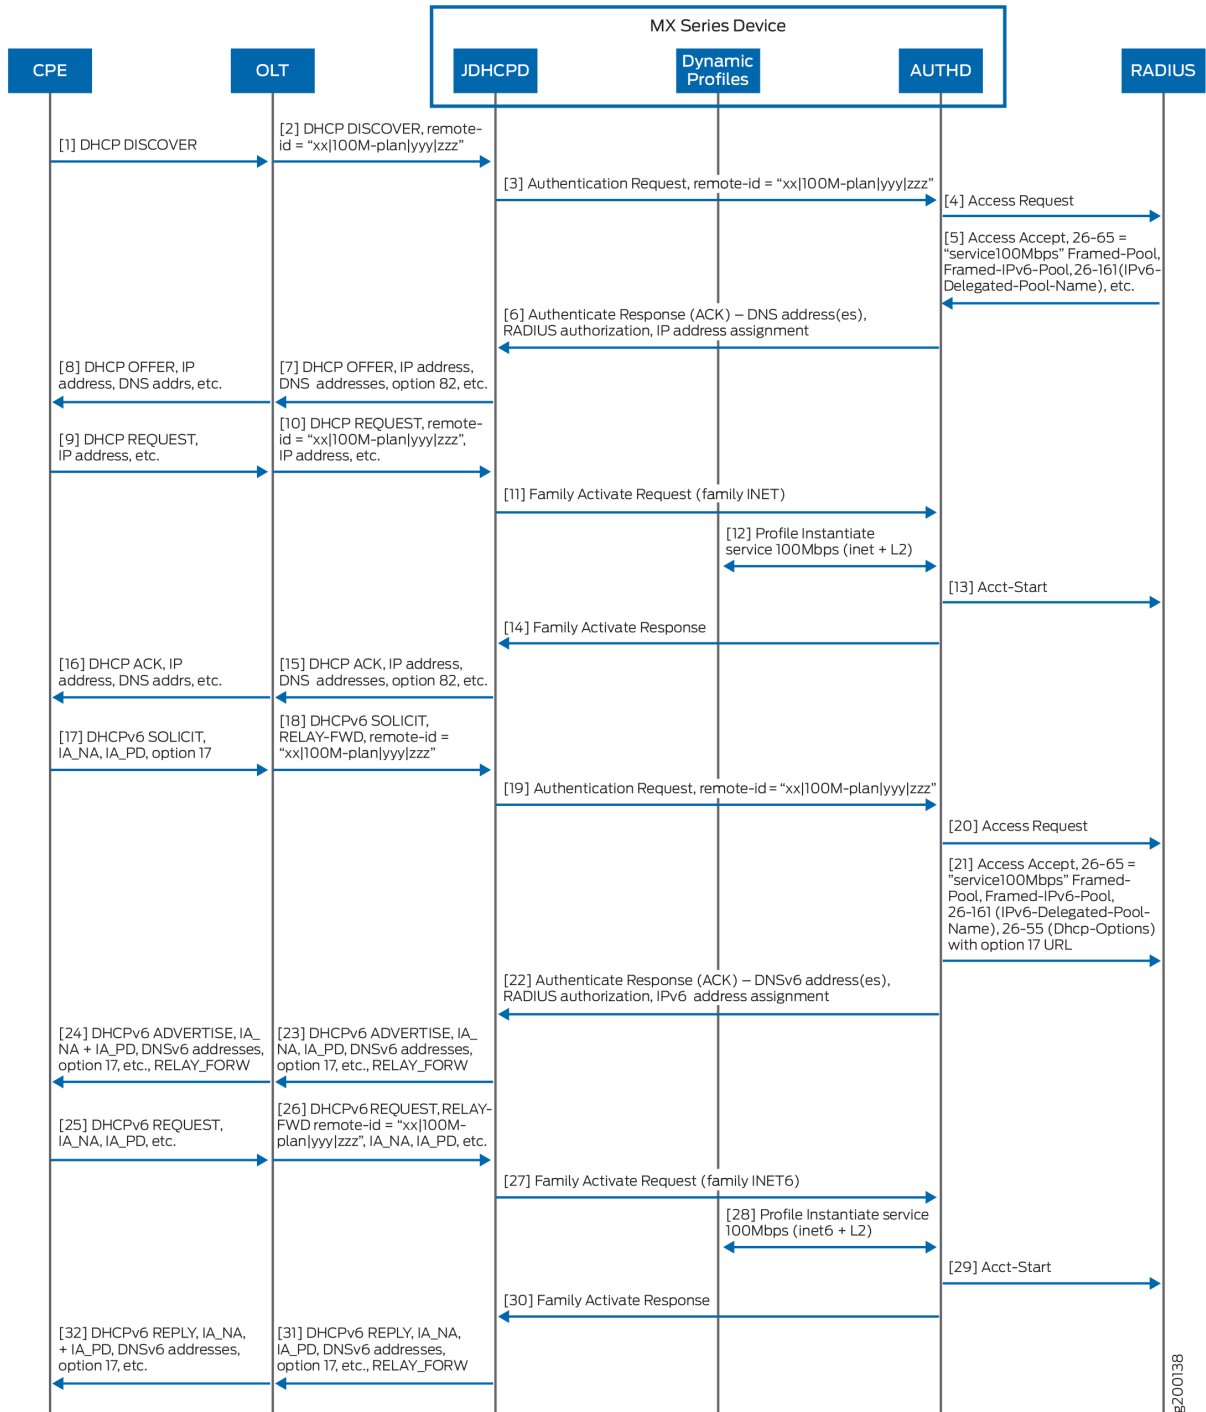
[Figure 1 on page 734](#) illustrates the sequence of steps in the initial negotiation between a DHCP client, a DHCP server, and the RADIUS server. The following terms are used in the figure:

CPE—Customer premises equipment (functions as the DHCP client or subscriber).

OLT—Optical line terminator—for example, a DSL access multiplexer (DSLAM) or other aggregation device.

MX Series device—Functions as the DHCP server.


Figure 1: Initial Negotiation



Service Plan Change

Figure 2 on page 735 illustrates the sequence of steps in a change of service plans, from a 100 Mbps plan to a 1 Gbps plan.:

Figure 2: Service Plan



RADIUS Attributes Supported for Reauthentication

Table 12 on page 736 lists the RADIUS standard attributes and VSAs that can be processed during reauthentication when received in the RADIUS Access-Accept message, and describes how authd handle changes in attributes. Attribute processing is consistent with CoA request processing. The characteristics of the reauthenticating subscriber session change only if new values or new attributes are received in the Access-Accept message.

Table 12: RADIUS attributes supported by reauthentication

Attribute Number	Attribute Name	Result of Processing
8	Framed-IP-Address	A new value is stored in the subscriber session database; old data is overwritten.
22	Framed-Route	A new value is stored in the subscriber session database; old data is appended.
24	State	A new value is stored in the subscriber session database; old data is overwritten.
25	Class	A new value is stored in the subscriber session database; old data is overwritten.
26-4	Primary-DNS	A new value is stored in the subscriber session database; old data is overwritten.
26-5	Secondary-DNS	A new value is stored in the subscriber session database; old data is overwritten.
26-6	Primary-WINS	A new value is stored in the subscriber session database; old data is overwritten.
26-7	Secondary-WINS	A new value is stored in the subscriber session database; old data is overwritten.
26-55	DHCP-Options	Value is sent to the DHCP local server for processing the changes to the subscriber's DHCP configuration.

Table 12: RADIUS attributes supported by reauthentication (Continued)

Attribute Number	Attribute Name	Result of Processing
26-65	Activate-Service	<p>The authd process compares the list of services in the VSA to the services that are already active for that subscriber session.</p> <ul style="list-style-type: none"> • If the list on the VSA contains services not yet active, authd activates those services for the subscriber. • If any service already active for the subscriber session is not listed in the VSA, then authd deactivates that service. <p>For example, suppose services A and B are active on the session, but the VSA includes only services B and C. Service A is not on the VSA list and is deactivated. Service C is on the list but not currently active, so authd activates C. Service B is both already active and on the list, so it remains active.</p>
26-161	IPv6-Delegated-Pool-Name	A new value is stored in the subscriber session database; old data is overwritten.
26-206	Reauthenticate-On-Renew	<p>If the value is 1 (enable), authd adds the value to the subscriber session database if it is not already present.</p> <p>If the value is 0 (disable) and a value of 1 is already present in the database, authd sets the database value to 0.</p> <p>If the value in the message is missing or invalid, and a value is already present in the database, authd deletes the value from the database.</p>
26-207	DHCPv6-Options	Value is sent to the DHCPv6 local server for processing the changes to the subscriber's DHCP configuration.
88	Framed-Pool	A new value is stored in the subscriber session database; old data is overwritten.
97	Framed-IPv6-Prefix	A new value is stored in the subscriber session database; old data is overwritten.

Table 12: RADIUS attributes supported by reauthentication (Continued)

Attribute Number	Attribute Name	Result of Processing
100	Framed-IPv6-Pool	A new value is stored in the subscriber session database; old data is overwritten.
123	Delegated-IPv6-Prefix	A new value is stored in the subscriber session database; old data is overwritten.
168	Framed-IPv6-Address	A new value is stored in the subscriber session database; old data is overwritten.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, when the DHCPv6 server discovers an issue with the address, a status code for NotOnLink is supported in addition to the previously supported codes for NoAddrsAvail and NoPrefixAvail.
18.1R1	The discover and solicit messages support reauthentication starting in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

[Configuring RADIUS Reauthentication for DHCP Subscribers | 738](#)

[Single-Session DHCP Dual-Stack Overview | 1316](#)

[Understanding Differences Between Legacy DHCP and Extended DHCP | 1015](#)

Configuring RADIUS Reauthentication for DHCP Subscribers

You can configure reauthentication as an alternative to RADIUS CoA messages as a means to change characteristics of the subscriber session, such as activating or changing service plans or changing DHCP subscriber attributes. When configured, reauthentication is triggered when the DHCP local server

receives a renew, rebind, discover, or solicit message from a DHCP client. The message triggers `jdhcpd` to request reauthentication from `authd`, which in turn reissues the RADIUS Access-Request for a second subscriber authentication. Reauthentication is available for DHCPv4, DHCPv6, and dual-stack subscribers.

Starting in Junos OS Release 18.1R1, reauthentication can be triggered by discover and solicit messages in addition to the previously supported renew and rebind messages. The release also introduces reauthentication support for dual-stack, single-session subscribers.

You can use the `reauthenticate` statement to configure reauthentication to occur in response to all DHCP renew, rebind, discover, or solicit messages or only in response to those messages when they include a different Agent Remote ID for the DHCP client. The Agent Remote ID carries information about the subscriber's service plan, so a change in ID value corresponds to a change in the subscriber service plan. The Agent Remote ID is conveyed in option 82, suboption 2 for DHCPv4 clients and in option 37 for DHCPv6 clients.

You can also use the Juniper Networks VSA, Reauthentication-On-Renew (26-206) as an alternative to the CLI configuration to enable reauthentication. The VSA is conveyed in the RADIUS Access-Accept message at subscriber login, and must be configured on your RADIUS server. The `reauthenticate` statement overrides the VSA when the VSA is present with a value of `disable`.

Configure reauthentication for non-dual-stack, single session DHCP subscribers:

- (Optional) Specify reauthentication is triggered by receipt of every renew, rebind, discover, and solicit message.

For DHCPv4 subscribers:

```
[edit system services dhcp-local-server]
user@host# set reauthenticate lease-renewal
```

For DHCPv6 subscribers:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set reauthenticate lease-renewal
```

- (Optional) Specify reauthentication is triggered only when the Agent Remote ID has changed in the received discover or solicit message.

For DHCPv4 subscribers:

```
[edit system services dhcp-local-server]
user@host# set reauthenticate remote-id-mismatch
```

For DHCPv6 subscribers:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set reauthenticate remote-id-mismatch
```

Configure reauthentication for dual-stack, single session DHCP subscribers:

1. Configure addresses to be allocated on demand for subscribers in the dual-stack group.

```
[edit system services dhcp-local-server dual-stack-group name]
user@host# set on-demand-address-allocation
```

2. (Optional) Specify reauthentication is triggered for every subscriber in the dual-stack group by receipt of every renew, rebind, discover, and solicit message.

```
[edit system services dhcp-local-server dual-stack-group name]
user@host# set reauthenticate lease-renewal
```

3. (Optional) Specify reauthentication is triggered for every subscriber in the dual-stack group only when the Agent Remote ID has changed in the received discover or solicit message.

```
[edit system services dhcp-local-server dual-stack-group name]
user@host# set reauthenticate remote-id-mismatch
```

A change in the Agent Remote ID can also initiate a service change during renew and rebind operations when the `remote-id-mismatch` statement is configured. You cannot configure both the `remote-id-mismatch` statement and the `reauthenticate` statement at the global level, `[edit system services dhcp-local-server]`. However, DHCP precedence rules do permit you to configure both statements when they are at different levels. For example, you can configure `reauthenticate` at the global level and `remote-id-mismatch` for DHCPv6 at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level or for a specific group at the `[edit system services dhcp-local-server group name]` hierarchy level, and so on.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, reauthentication can be triggered by discover and solicit messages in addition to the previously supported renew and rebind messages. The release also introduces reauthentication support for dual-stack, single-session subscribers.

RELATED DOCUMENTATION

[RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 727](#)

[Single-Session DHCP Dual-Stack Overview | 1316](#)

[DHCP-Initiated Service Change Based on Remote ID | 1071](#)

[Understanding Differences Between Legacy DHCP and Extended DHCP | 1015](#)

RADIUS Accounting for Subscriber Access

IN THIS SECTION

- [RADIUS Accounting Statistics for Subscriber Access Overview | 742](#)
- [RADIUS Acct-On and Acct-Off Messages | 744](#)
- [Configuring Per-Subscriber Session Accounting | 744](#)
- [Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI | 748](#)
- [Understanding RADIUS Accounting Duplicate Reporting | 749](#)
- [Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting | 751](#)
- [Configuring Per-Service Session Accounting | 752](#)
- [Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 754](#)
- [Configuring Service Packet Counting for Volume Statistics | 756](#)
- [Configuring Service Accounting | 757](#)
- [Preservation of RADIUS Accounting Information During an Accounting Server Outage | 759](#)
- [Configuring Back-up Options for RADIUS Accounting | 762](#)
- [Forcing the Router to Contact the Accounting Server Immediately | 763](#)

- [Monitoring Pending RADIUS Accounting Stop Messages | 764](#)
- [Suspending RADIUS Accounting and Baseline Accounting Statistics Overview | 766](#)
- [Configuring RADIUS Accounting Suspension and Baseline Accounting Statistics | 770](#)

This topic provides detailed information about RADIUS accounting statistics, subscriber session accounting, duplicate reporting, and service accounting. For information about configuring servers for RADIUS accounting, see ["RADIUS Authentication and Accounting Basic Configuration" on page 721](#).

RADIUS Accounting Statistics for Subscriber Access Overview

The AAA Service Framework enables you to configure how the router collects and uses accounting statistics for subscriber management.

For example, you can specify when statistics collection is terminated, the order in which different accounting methods are used, the types of statistics collected, and how often statistics are collected. You can also configure the router to request that the RADIUS server immediately update the accounting statistics when certain events occur, such as when a subscriber logs in or when a change of authorization (CoA) occurs.

Subscriber management provides two levels of subscriber accounting—subscriber session and service session. In subscriber session accounting, the router collects statistics for the entire subscriber session. In service session accounting, the router collects statistics for specific service sessions for the subscriber.



NOTE: Subscriber management counts forwarded packets only. Dropped traffic (for example, as a result of a filter action) and control traffic are not included in the accounting statistics.

The router uses the RADIUS attributes and Juniper Networks VSAs listed in [Table 13 on page 743](#) to provide the accounting statistics for subscriber and service sessions. If the session has both IPv4 and IPv6 families enabled, the router reports statistics for both families.



NOTE: RADIUS reports subscriber statistics as an aggregate of both IPv4 statistics and IPv6 statistics.

- For an IPv4-only configuration, the standard RADIUS attributes report the IPv4 statistics and the IPv6 VSA results are all reported as 0.

- For an IPv6-only configuration, the standard RADIUS attributes and the IPv6 VSA statistics are identical, both reporting the IPv6 statistics.
- When both IPv4 and IPv6 are configured, the standard RADIUS attributes report the combined IPv4 and IPv6 statistics. The IPv6 VSAs report IPv6 statistics.

Table 13: RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting

Attribute Number	Attribute Name	Type of Statistics
26-151	IPv6-Acct-Input-Octets	IPv6
26-152	IPv6-Acct-Output-Octets	IPv6
26-153	IPv6-Acct-Input-Packets	IPv6
26-154	IPv6-Acct-Output-Packets	IPv6
26-155	IPv6-Acct-Input-Gigawords	IPv6
26-156	IPv6-Acct-Output-Gigawords	IPv6
47	Acct-Input-Packets	IPv4 and IPv6 aggregation
48	Acct-Output-Packets	IPv4 and IPv6 aggregation
52	Acct-Input-Gigawords	IPv4 and IPv6 aggregation
53	Acct-Output-Gigawords	IPv4 and IPv6 aggregation

SEE ALSO

[RADIUS Authentication and Accounting Basic Configuration](#) | 721

RADIUS Acct-On and Acct-Off Messages

Subscriber management supports RADIUS Acct-On and Acct-Off messages to indicate the current state of RADIUS accounting support.

RADIUS Acct-On messages indicate that accounting is being supported. Subscriber management issues Acct-On messages in the following situations:

- Accounting is enabled through configuration (for example, an accounting server is configured).
- A new access profile is configured and committed for a logical system/routing instance context. However, no Acct-On message is sent if the accounting server exists prior to the access profile and if it is simply modified.
- The router performs a cold reboot.
- The router performs a warm reboot and there are no subscribers currently logged in.
- The Authd process restarts and there are no active subscribers.

RADIUS Acct-Off messages indicate that accounting is not supported. Subscriber management issues Acct-Off messages in the following situations:

- The Authd process is terminated and there are no active subscribers.
- The router is shut down and accounting servers are currently configured (this action also logs out all current subscribers).
- The router is rebooted and redundancy is disabled.

SEE ALSO

[AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs](#) | 614

Configuring Per-Subscriber Session Accounting

To configure accounting for a subscriber session, you use an access profile, and specify how the subscriber access management feature collects and uses the accounting statistics. The router uses the RADIUS attributes and Juniper Networks VSAs discussed in "[RADIUS Accounting Statistics for Subscriber Access Overview](#)" on page 742 to provide the accounting statistics for the subscriber session.

To configure accounting for a subscriber session:

1. At the [edit access profile *profile-name*] hierarchy level, specify that you want to configure accounting.

```
[edit access profile profile-name]  
user@host# edit accounting
```

2. (Optional) Configure AAA to issue an Acct-Stop message if the AAA server denies access to the subscriber.

```
[edit access profile profile-name accounting]  
user@host# set accounting-stop-on-access-deny
```

3. (Optional) Configure AAA to send an Acct-Stop message if the subscriber fails AAA but is granted access by the AAA server.

```
[edit access profile profile-name accounting]  
user@host# set accounting-stop-on-failure
```

4. (Optional) Configure the router or switch to send an Acct-Update message to the RADIUS accounting server when a CoA occurs.

```
[edit access profile profile-name accounting]  
user@host# set coa-immediate-update
```

5. (Optional) Configure subscriber management to send the RADIUS accounting report to both the wholesaler and the retailer accounting servers.

```
[edit access profile profile-name accounting]  
user@host# set duplication
```

6. (Optional) Configure the duplication filtering action you want the router to perform when the RADIUS duplication accounting operation is enabled.

```
[edit access profile profile-name accounting]  
user@host# set duplication-filter interim-duplicated exclude-attributes
```


7. (Optional) Configure the router to send the RADIUS accounting report to multiple accounting servers listed in access profiles in a nondefault VRF (LS:RI).

```
[edit access profile profile-name accounting duplication-vrf]
user@host# set vrf-name vrf-name
user@host# set access-profile-name profile-name
```

8. (Optional) Configure the router or switch to send an Acct-Update message to the RADIUS accounting server when the router or switch receives a response (for example, an ACK or timeout) to the Acct-Start message.

```
[edit access profile profile-name accounting]
user@host# set immediate-update
```

9. (Optional) Configure the order in which multiple accounting methods are used.

```
[edit access profile profile-name accounting]
user@host# set order [ accounting-order ]
```

10. (Optional) Configure the types of statistics to gather. You can specify that the router or switch collect both volume and time statistics or only time statistics for subscriber sessions. When you change the type of statistics being collected, current subscribers continue to use the previous collection specification. Subscribers who log in after the change use the new specification.

```
[edit access profile profile-name accounting]
user@host# set statistics (time | volume-time)
```

11. (Optional) Override the default behavior and specify that, after a CoA action that changes the RADIUS Class attribute, accounting reports for the subscriber's service sessions continue to use the original Class attribute that was assigned when the service sessions were created. The new Class attribute value is used in accounting reports for the subscriber session only. By default, the accounting reports for both the subscriber session and the subscriber's service sessions use the new Class attribute value.

```
[edit access profile profile-name accounting]
user@host# set coa-no-override service-class-attribute
```

12. (Optional) Configure the number of minutes between accounting updates. You can configure an interval from 10 through 1440 minutes. All values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

```
[edit access profile profile-name accounting]
user@host# set update-interval minutes
```

13. (Optional) Configure AAA to send an immediate interim accounting update to the RADIUS server when AAA receives a rate change notification from the ANCP agent on the router.

```
[edit access profile profile-name accounting]
user@host# set ancp-speed-change-immediate-update
```

14. (Optional) Configure the authd process to wait for an Acct-On-Ack response message from RADIUS before sending any new authentication and accounting updates to the RADIUS server. This configuration ensures that when a new subscriber session starts, the authentication and accounting information for the new session does not get deleted when RADIUS clears previously existing session state information.

```
[edit access profile profile-name accounting]
user@host# set wait-for-acct-on-ack
```

15. (Optional) Configure the authd process to send accounting messages when the RADIUS server status changes for an access profile. It sends an Acct-On message when the first RADIUS server is added to the access profile and sends an Acct-Off message when the last RADIUS server is removed from the access profile. This configuration enables you to monitor whether the access profile has an active RADIUS server.

```
[edit access profile profile-name accounting]
user@host# set send-acct-status-on-config-change
```

SEE ALSO

[Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | 1714](#)

[RADIUS Authentication and Accounting Basic Configuration | 721](#)

Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI

You can configure the router to display accurate statistics for subscriber sessions on dynamic interfaces. By default, aggregate statistics (byte and packet counts) for interfaces displayed by the `show interfaces extensive` command do not accurately reflect customer traffic. These counters include overhead bytes that represent the encapsulation overhead added to the actual subscriber data bytes. The aggregate counters also include dropped packets in the total, so the values represent transit statistics rather than the actual subscriber traffic on the interface.

Inclusion of the overhead bytes and dropped packets can have a significant effect on the final reported values. You can exclude dropped packets from the count by including the `interface-transmit-statistics` statement for an interface, but this has no effect on the overhead bytes.

To display accurate subscriber statistics, include the `actual-transmit-statistics` statement for the logical interface in the dynamic profile. This statement enables the `show subscribers` command to display aggregate byte and packet counts for a specified subscriber session or for all subscriber sessions on a specified interface. The displayed statistics match the values that are reported to RADIUS for the subscribers. The statistics are collected after traffic shaping is applied and they do not include overhead bytes, control packets, or dropped packets.



NOTE: Starting in Junos OS Release 18.4R1, you must enable `actual-transit-statistics` to collect subscriber statistics. If you do not configure this statement, subscriber statistics are not collected; the `show subscribers accounting-statistics` command displays a value of 0 for subscriber statistics; and the subscriber statistics are reported to RADIUS with values of zero.



NOTE: Service accounting statistics are not included.

To configure the reporting of accurate subscriber session statistics:

- Enable actual transit statistics.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-
interface-unit]
user@host# set actual-transit-statistics
```

You can display the subscriber accounting statistics in two ways:

- Display subscriber statistics by session ID with the `show subscribers id session-id accounting-statistics` command.

- Display subscriber statistics by dynamic interface for all session IDs with the `show subscribers interfaces interface-name accounting-statistics` command.

Understanding RADIUS Accounting Duplicate Reporting

IN THIS SECTION

- [Layer 3 Wholesale Scenarios | 749](#)
- [Other Scenarios | 750](#)
- [Filters for Duplicate Accounting Reports | 751](#)

When you configure RADIUS accounting, by default the router sends the accounting reports to the accounting servers in the context in which the subscriber was last authenticated. You can configure RADIUS accounting to send duplicate accounting reports to other servers in the same context or in other contexts.

Layer 3 Wholesale Scenarios

In a Layer 3 wholesale network environment, the wholesaler and retailer might use different RADIUS accounting servers, and both might want to receive accounting reports. In this situation, you can configure RADIUS accounting duplicate reporting, which sends reports to both the wholesaler and the retailer accounting servers. The target to which the duplicate accounting records are sent must be in the default:default logical system:routing instance combination (LS:RI) , also called the *default VRF*.

[Table 14 on page 750](#) shows where subscriber management sends the accounting reports when you enable duplicate reporting. Subscriber management sends duplicate reports based on the access profile in which you configure the duplication statement at the [edit access profile *profile-name* accounting] hierarchy level, where the subscriber resides, and how the subscriber is authenticated.



NOTE: You can also enable accounting duplicate reporting based on the domain map configuration—you configure subscribers to authenticate with a nondefault routing instance and a target logical system:routing instance of default:default. The accounting reports are then sent to both the authentication context and the default:default context.

Table 14: Duplicate RADIUS Accounting Reporting

Access Profile in Which Duplication Is Configured	Where Subscriber Is Authenticated	Subscriber's Target Logical System/Routing Instance	Accounting Servers Where Accounting Reports Are Sent
retailer A	wholesaler	retailer A	wholesaler and retailer A
retailer A	retailer A	retailer A	wholesaler (default/default context) NOTE: This is the domain map configuration described in the Note preceding this table.
wholesaler	wholesaler and retailer A	retailer A	wholesaler and retailer A
wholesaler and retailer B	wholesaler and retailer A	retailer B	wholesaler, retailer A, and retailer B
not configured (default)	any	any	single report sent to accounting servers in the context in which subscriber was last authenticated

Other Scenarios

For scenarios that are not in a Layer 3 wholesale network environment, you might want to send duplicate accounting records to a different set of RADIUS servers that reside in either the same or a different routing context. Unlike the Layer 3 wholesale scenario, the target for the duplicate RADIUS accounting records does not have to be the default VRF. You can specify a single nondefault VRF—that is, other than the default:default LS:RI combination—as the target. Additionally, you can specify up to five access profiles in the target VRF that list the RADIUS accounting servers that receive the duplicate reports.

For example, you might have a lawful intercept scenario where the subscriber is authenticated in the default domain. An authorized law enforcement organization needs duplicate accounting records for the

subscriber to be sent to a mediation device that resides in the organization's networking domain, which lies in a nondefault VRF.

Subscriber management sends duplicate reports to the VRF that you specify with the `vrf-name` statement at the `[edit access profile profile-name accounting duplication-vrf]` hierarchy level. Include the `access-profile-name` statement at the same level to designate the access profiles that in turn specify the RADIUS servers that receive the duplicate reports.

Filters for Duplicate Accounting Reports

Subscriber management provides a duplication filter feature that enables you to specify which accounting servers receive the RADIUS accounting interim reports when RADIUS accounting duplicate reporting is active. You configure the filters in the AAA access profile, and the router then applies the filters to subscribers associated with that profile.

Subscriber management supports the following filtering for RADIUS accounting duplicate reporting:

- **Duplicated accounting interim messages**—The router filters duplicate accounting messages. The accounting messages are sent only to RADIUS accounting servers in the subscriber's access profile.
- **Original accounting interim messages**—The router filters accounting messages destined for original RADIUS accounting servers, which are accounting servers in the subscriber's access profile. The accounting messages are sent only to duplication accounting servers (servers in a duplication access profile other than the subscriber's access profile).
- **Excluded RADIUS attributes**—The router filters the RADIUS attributes in the accounting messages based on the `exclude` statement configuration in the access profile under the duplication context. You can use the `exclude` filter alone, or with the duplicated or original accounting message filters.

Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting

You can use duplication filters to specify the RADIUS accounting servers that receive RADIUS accounting interim reports when accounting duplicate reporting is enabled. You configure the filters in a AAA access profile, and the router applies the filters to subscribers associated with that profile.

To configure duplication filters for accounting duplicate reporting:

1. At the `[edit access profile profile-name]` hierarchy level, specify that you want to configure accounting.

```
[edit access profile profile-name]  
user@host# edit accounting
```

2. Configure the duplication filter you want the router to use.

The following examples show the three types of filters and describe the results for each filter:

- Specify that the router does not send the accounting interim messages to duplicate RADIUS accounting servers.

Duplicate RADIUS accounting servers are those that are not in the subscriber's access profile. The router still sends the accounting interim messages to accounting servers that reside in the subscriber's access profile.

```
[edit access profile profile-name accounting]
user@host# set duplication-filter interim-duplicated
```

- Specify that the router does not send the accounting interim messages to original RADIUS accounting servers.

Original accounting servers are those that reside in the subscriber's AAA routing context. The router still sends the accounting interim messages to duplicate accounting servers, which are those servers that do not reside in a duplication context other than the subscriber's access profile.

```
[edit access profile profile-name accounting]
user@host# set duplication-filter interim-original
```

- Specify how the router uses the `exclude` statement configuration to filter RADIUS attributes from accounting interim messages.

The router uses the configuration for the `exclude` statement in the duplication access profile to determine which RADIUS attributes are not included in the accounting interim messages.

```
[edit access profile profile-name accounting]
user@host# set duplication-filter exclude-attributes
```

Configuring Per-Service Session Accounting

Subscriber management enables you to configure the router to collect statistics on a per-service session basis for subscribers. Per-service session accounting requires two operations. First, RADIUS must be configured to provide the name of the service, the accounting interval to use, and the type of statistics to collect (either time statistics or a combination of time and volume statistics). Second, if RADIUS VSA 26-69 is configured for time and volume statistics, you must also configure a firewall or fast update firewall filter that counts service packets—the service packet information provides the volume statistics.

The router uses the RADIUS attributes and Juniper Networks VSAs discussed in ["RADIUS Accounting Statistics for Subscriber Access Overview" on page 742](#) to provide the accounting statistics for the subscriber session.



NOTE: The collection of time-only service statistics is supported for all service sessions. However, time and volume statistics are provided for only firewall and fast update firewall service sessions.

To configure the router to provide per-service accounting statistics:

1. Ensure that the required RADIUS VSAs are configured.
See [Table 15 on page 753](#) for the VSAs that the router uses for per-service accounting.
2. Configure the classic firewall filter or fast update filter to count the service packets.
See ["Configuring Service Packet Counting for Volume Statistics" on page 756](#).

Table 15: Juniper Networks VSAs Used for Per-Service Session Accounting

Attribute Number	Attribute Name	Description	Value
26-69	Service-Statistics	Enable or disable statistics for the service	<ul style="list-style-type: none"> 0 = disable 1 = enable time statistics 2 = enable time and volume statistics
26-83	Service-Session	Service string sent in accounting stop and start messages from the router to the RADIUS server	string: service-name, with parameter values that are sent from RADIUS server in attribute 26-65.
26-140	Service-Interim-Acct-Interval	Amount of time between interim accounting updates for this service	<ul style="list-style-type: none"> range = 600–86400 seconds 0 = disabled <p>NOTE: Values are rounded up to the next higher multiple of 10 minutes. For example, a setting of 900 seconds (15 minutes) is rounded up to 20 minutes (1200 seconds).</p>

SEE ALSO

[RADIUS Authentication and Accounting Basic Configuration](#) | 721

Processing Cisco VSAs in RADIUS Messages for Service Provisioning

You can use Cisco VSAs in RADIUS messages to provision and manage services in a subscriber access network. In the topology for this deployment, the broadband network gateway (BNG) is connected to:

- A RADIUS server, such as the Steel-Belted Radius Carrier (SBRC), that is used to authentication and accounting.
- A Cisco BroadHop application that is used as the Policy Control and Charging Rules Function (PCRF) server for provisioning services using RADIUS change of authorization (CoA) messages.

Cisco BroadHop does not support Juniper VSAs. It uses the Cisco VSA, Cisco-AVPair (26-1, IANA private enterprise number 9) with different values to activate and deactivate the services.

To activate a service, use the Cisco-AVPair VSA (26-1) with each of the following values:

- Value of the *.subscriber:command=activate-service* parameter.
- Value of the *subscriber:service-name=service-name* parameter.

To deactivate a service, use the Cisco-AVPair VSA (26-1) with each of the following values:

- Value of the *subscriber:command=deactivate-service* parameter.
- Value of the *subscriber:service-name=service-name* parameter.

You cannot modify any attributes in authentication, accounting, or CoA responses in the RADIUS messages that the BNG sends. Any Cisco VSAs other than the ones used to provision the services are considered as unsupported attributes.

To configure service accounting for an access profile for a subscriber:

1. Specify that you want to configure service accounting.

```
[edit access profile profile-name service]
user@host# edit accounting
```

2. (Optional) Enable interim service accounting updates and configure the amount of time that the router or switch waits before sending a new service accounting update. You can configure an interval

from 10 through 1440 minutes. All values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

```
[edit access profile profile-name service accounting]
user@host# set update-interval minutes
```

3. (Optional) Configure the types of statistics to gather. You can specify that the router or switch collect both volume and time statistics or only time statistics for subscriber sessions. When you change the type of statistics being collected, current subscribers continue to use the previous collection specification. Subscribers who log in after the change use the new specification.

```
[edit access profile profile-name service accounting]
user@host# set statistics (time | volume-time)
```

You can also define the UDP port number to configure the port on which the router that functions as the RADIUS dynamic-request server must receive requests from RADIUS servers. By default, the router listens on UDP port 3799 for dynamic requests from remote RADIUS servers. You can configure the UDP port number to be used for dynamic requests for a specific access profile or for all of the access profiles on the router. To define the UDP port number, include the `dynamic-request-port port-number` statement at the `[edit access profile profile-name radius-server server-address]` or the `[edit access radius-server server-address]` hierarchy level.

To specify the UDP port globally for all access profiles:

```
[edit access radius-server server-address]
user@host# set dynamic-request-port port-number
```

To specify the UDP port for a specific access profile:

```
[edit access profile profile-name radius-server server-address]
user@host# set dynamic-request-port port-number
```

SEE ALSO

| [Standard and Vendor-Specific RADIUS Attributes](#) | 3

Configuring Service Packet Counting for Volume Statistics

Subscriber management uses service packet counting to report volume statistics for subscribers on a per-service session basis. To configure service packet counting, you specify the accounting action, and subscriber management then applies the results to a specific named counter (`__junos-dyn-service-counter`) for use by RADIUS.

The accounting action you configure specifies the counting mechanism that subscriber management uses when capturing statistics—either inline counters or deferred counters. Inline counters are captured when the event occurs, and do not include any additional packet processing that might occur after the event. Deferred counters (also called accurate accounting) are not incremented until the packet is queued for transmission, and therefore include the entire packet processing. Deferred counters provide a more accurate count of the packets than inline counters, and are more useful for subscriber accounting and billing.

You configure the accounting mechanism by specifying either the `service-accounting-deferred` action (for deferred counters) or the `service-accounting` action (for inline counters) at the `[edit firewall family family-name filter filter-name term term-name then] hierarchy level`.

The two accounting mechanisms are mutually exclusive, both on a per-term basis and a per-filter basis. Also, both accounting actions are mutually exclusive with the `count` action on a per-term basis.



NOTE: You can define deferred counters for the `inet` and `inet6` families for classic filters only. Fast update filters do not support deferred counters.

To enable service packet counting:

1. Configure any match conditions that you want to count using the service accounting action. For example:

```
[edit firewall family inet filter filtername term term-name]
user@host# set from source-address address
```

2. Specify the accounting action for the filter.

To use deferred counters:

```
[edit firewall family inet filter filtername term term-name]
user@host# set then service-accounting-deferred
```

To use inline counters:

```
[edit firewall family inet filter filtername term term-name]
user@host# set then service-accounting
```

When the match conditions for the filter are met, the packet is counted and applied to the well-known service counter (`__junos-dyn-service-counter`) for use by the RADIUS server. This counter provides the volume statistics for per-service accounting.



TIP: You cannot use the `service-accounting` action or the `service-accounting-deferred` action in the same term as a `count` action.

SEE ALSO

Classic Filters Overview

Defining Dynamic Filter Processing Order

[Guidelines for Configuring Firewall Filters](#)

[Guidelines for Applying Standard Firewall Filters](#)

[Firewall Filter Terminating Actions](#)

[Firewall Filter Nonterminating Actions](#)

Configuring Service Accounting

Service accounting is disabled by default. You can configure service accounting by using RADIUS attributes received from the external RADIUS server or by using the CLI to configure accounting locally on the router. If you configure both, the RADIUS setting takes precedence over the CLI setting.

In some networks, you must use the CLI to enable and disable service accounting and to specify the interim accounting interval. For example, the BNG might be connected to both a RADIUS server and a third-party device using an application that uses RADIUS CoAs for service provisioning but does not support Juniper Networks VSAs. For more information about this use case, see ["Processing Cisco VSAs in RADIUS Messages for Service Provisioning"](#) on page 754.

[Table 16 on page 758](#) indicates the type of service accounting statistics that are collected when various combinations of local CLI and RADIUS service accounting configuration are present:

Table 16: Type of Service Accounting Statistics Collected Based On CLI and RADIUS Configurations

CLI Configuration Present for Service Statistics	RADIUS Configuration Present for Service Statistics	Service Statistics Collected
–	–	None
–	✓	RADIUS configuration
✓	–	CLI configuration
✓	✓	RADIUS configuration
✓	Explicitly disabled with a value of 0	None

Table 17 on page 758 indicates the service interim accounting interval value that is used when various combinations of local CLI and RADIUS service accounting configuration are present:

Table 17: Service Interim Accounting Interval Value Based on CLI and RADIUS Configurations

CLI Configuration Present for Service Interim Accounting Interval	RADIUS Configuration Present for Service Interim Accounting Interval	Service Interim Accounting Interval Value Used
–	–	No service interim accounting
–	✓	RADIUS value
✓	–	CLI value
✓	✓	RADIUS value
✓	Explicitly disabled with a value of 0	No service interim accounting

Table 18 on page 759 shows the results for two example combinations of CLI and RADIUS configurations.

Table 18: Example of Values Used for Different Configurations

CLI	RADIUS	Value Used
update-interval = 400	Acct-Interim-Interval (85) = 600	600
statistics = time	Service-Statistics (26-69) not set	time
update-interval = 400	Acct-Interim-Interval (85) not set	400
statistics = time	Service-Statistics (26-69) = 2, time and volume	time and volume

To configure service accounting for an access profile for a subscriber:

1. Specify that you want to configure service accounting.

```
[edit access profile profile-name service]
user@host# edit accounting
```

2. (Optional) Enable interim service accounting updates and configure the amount of time that the router or switch waits before sending a new service accounting update. You can configure an interval from 10 through 1440 minutes. All values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

```
[edit access profile profile-name service accounting]
user@host# set update-interval minutes
```

3. (Optional) Configure the types of statistics to gather. You can specify that the router or switch collect both volume and time statistics or only time statistics for subscriber sessions. When you change the type of statistics being collected, current subscribers continue to use the previous collection specification. Subscribers who log in after the change use the new specification.

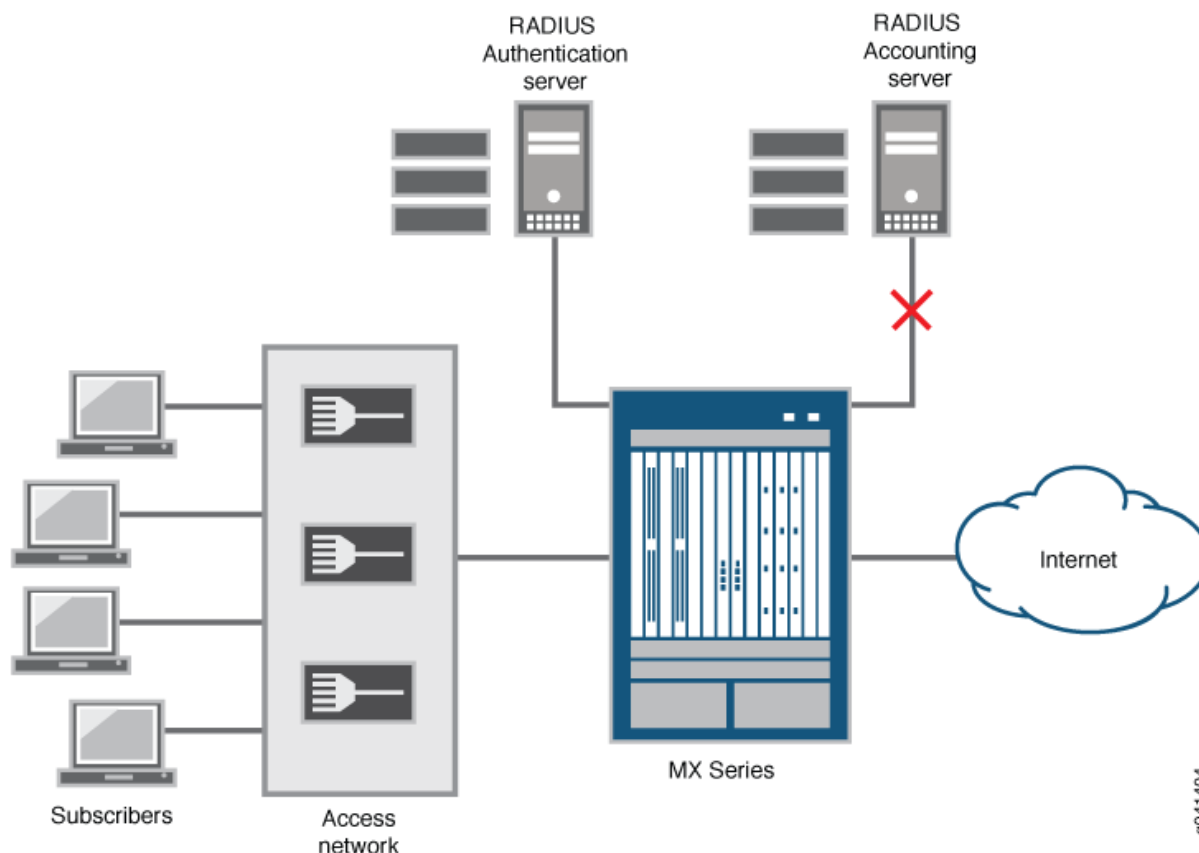
```
[edit access profile profile-name service accounting]
user@host# set statistics (time | volume-time)
```

Preservation of RADIUS Accounting Information During an Accounting Server Outage

If the router loses contact with the RADIUS accounting server, as represented in [Figure 3 on page 760](#), whether due to a server outage or a problem in the network connecting to the server, you can lose all

the billing information that would have been received by the server. RADIUS accounting backup preserves the accounting data that accumulates during the outage. If you have not configured RADIUS accounting backup, the accounting data is lost for the duration of the outage from the time when the router has exhausted its attempts to resume contact with the RADIUS server. The configurable retry value determines the number of times the router attempts to contact the server.

Figure 3: Topology with Loss of Access to Accounting Server



By default, the router must wait until the revert timer expires before it can attempt to contact the non-responsive server again. However, when you configure accounting backup, the revert timer is disabled and the router immediately retries its accounting requests as soon as the router fails to receive accounting acknowledgments. Accounting backup follows this sequence:

1. The router fails to receive accounting acknowledgments from the server.
2. The router immediately attempts to contact the accounting server and marks the server as offline if the router does not receive an acknowledgment before exhausting the number of retries.
3. The router next attempts to contact in turn each additional accounting server configured in the RADIUS profile.

If a server is reached, then the router resumes sending accounting requests to this server.

4. If none of the servers responds or if no other servers are in the profile, the router declares a timeout and begins backing up the accounting data. It withholds all accounting stop messages and does not forward new accounting requests to the server.
5. During the outage, the router sends a single pending accounting stop message to the servers at periodic intervals.
6. If one of the servers acknowledges receipt, then the router sends all the pending stop messages to that server in batches at the same interval until all the stored stop messages have been sent. However, any new accounting requests are sent immediately rather than being held and sent periodically.

The router replays accounting stop messages to the server in the correct order because it preserves both the temporal order among subscribers and the causal order between service and session stop requests for each subscriber. Only accounting stop messages are backed up, because they include the start time and duration of sessions and all the accounting statistics. This makes it unnecessary to withhold the accounting start messages, which eventually time out. Interim updates are not backed up and time out as well; if the session remains active, then the next interim update after the server connection is restored provides the interim accounting information.

You can configure the number of accounting stop messages that the router can queue pending restoration of contact with the accounting server. To preserve current accounting data in preference to collecting new accounting data, subscriber logins fail as soon as the maximum number of messages has been withheld. Subscriber logins resume immediately when the pending queue drops below the queue limit.



NOTE: Service accounting stop messages are withheld for a maximum of ten services per subscriber. If a subscriber attempts to activate an eleventh service while that accounting server is offline, the activation fails.

The router can hold the pending accounting messages for up to 24 hours. When the configurable maximum holding period passes, all accounting stop messages still in the pending queue are flushed, even if the accounting server has come back online. A consequence of this is that subscriber logins resume immediately if they were failing because the maximum pending limit had been reached.

All pending messages are also flushed in either of the following circumstances:

- If you remove the last accounting server from the access profile, because then there is no place to send the messages.
- If you remove the accounting backup configuration.

While the router is withholding accounting stop messages, you can force the router to attempt contact with the accounting server immediately, rather than allowing it to wait until the periodic interval has

expired. When you do so, the router first replays a batch of stop messages to the server, with one of the following outcomes:

- If the router receives an acknowledgment of receipt, then it marks the server as online and begins replaying all remaining pending stop messages in batches.
- If the router does not receive the acknowledgment, then it resumes sending a single pending accounting stop message at the periodic interval.

When a subscriber logs out while the accounting server is offline, the accounting stop requests for the subscriber and the session are queued and replayed to the server when it comes online. In this case, the subscriber session and service session information is retained, so that the router can send a correct accounting request when the server comes back online.

In the event of a *graceful Routing Engine switchover* while the accounting server is offline, the pending stop messages can be replayed from the active Routing Engine when the server is online again.



NOTE: When RADIUS accounting backup is configured, you must use different servers for RADIUS authentication and accounting. Subscriber authentication fails when the same server is configured for both authentication and accounting.

If the RADIUS server acts on behalf of other back-end RADIUS accounting or authentication servers and forwards requests to them, subscribers can be authenticated but accounting requests are not sent out.

Use the `show network-access aaa statistics` command to view backup accounting statistics.

Configuring Back-up Options for RADIUS Accounting

You can configure RADIUS accounting backup to preserve accounting data when the accounting server is unavailable because of a server or network outage. When backup is configured, RADIUS accounting stop messages are withheld and queued to be sent when connectivity is restored. You can specify the maximum number of stop messages that can be queued. When this maximum is reached, subsequent new subscriber logins fail because there is no remaining capacity to preserve accounting data for new sessions.

You can also configure how long the queued messages can be held. When this period expires, all pending accounting stops are flushed from the queue, even if the accounting server has come back online.



CAUTION: Before you configure RADIUS accounting backup, ensure that RADIUS accounting and RADIUS authentication are configured on different servers. Subscriber

authentication fails when the same server is configured for both authentication and accounting.

1. Enable accounting backup to use the default values.

```
[edit access ]
user@host# set accounting-backup-options
```

2. (Optional) Configure the number of accounting stops that the router can preserve while the accounting server is offline.

```
[edit access accounting-backup-options]
user@host# set max-pending-accounting-stops number
```

3. (Optional) Configure how long the router holds pending accounting stops before flushing them.

```
[edit access accounting-backup-options]
user@host# set max-withhold-time hold-time
```

For example, the following statements configure the backup options for all subscriber accounting; these statements specify that the router holds no more than 32,000 pending accounting stops—at which point all subsequent subscriber logins fail—and holds them no longer than 6 hours—at which point all pending messages are flushed and subscriber logins resume if they were failing:

```
[edit access accounting-backup-options]
user@host# set max-pending-accounting-stops 32000
user@host# set max-withhold-time 360
```

Use the `show network-access aaa statistics` command to view backup accounting statistics.

Forcing the Router to Contact the Accounting Server Immediately

In the event of an accounting server outage while RADIUS accounting backup is enabled, by default the router waits for a time interval to expire before contacting the offline server. Rather than waiting for that interval to pass, you can force the router to immediately contact the server by issuing the `request network-access aaa replay pending-accounting-stops` command. The router sends a batch of pending accounting stop requests to the server. If the router receives an acknowledgment from the server, then the router continues to replay the pending messages to the server in batches at the periodic interval. If

the router does not get that acknowledgment, then it resumes sending a single pending accounting stop message at the periodic interval.

To force the router to immediately contact the offline accounting server:

- Request the messages to be replayed.

```
user@host> request network-access aaa replay pending-accounting-stops
```

Monitoring Pending RADIUS Accounting Stop Messages

IN THIS SECTION

- Purpose | 764
- Action | 764

Purpose

Display information about RADIUS accounting stop messages that are being withheld due to an inability to contact the RADIUS accounting server.

Action

When you want to know whether the number of pending accounting-stop messages is nearing the maximum, you can display a simple count of pending requests:

```
user@host> show network-access aaa statistics pending-accounting-stops
Pending accounting stops: 10,000
```

You can use other commands to display more information about the accounting messages. The next example displays information for all services in the accounting session for the user, `vjshah29@example.com`. Although this example shows only one user, this command actually displays the information for all subscribers for whom accounting is being backed up.

```
user@host> show accounting pending-accounting-stops detail
Type: pppoe
Username: vjshah29@example.com
AAA Logical system/Routing instance: default:default
```

```

Access-profile: ce-ppp-profile
Session ID: 84
Accounting Session ID: 84
IP Address: 192.168.0.25
IPv6 Prefix: 2001:db8:2010:9999:18::/48
Authentication State: AuthAcctStopAckWait
Accounting State: Acc-Stop-Stats-Pending
Service name: cos-service
  Service State: SvcInactive
  Session ID: 94
  Session uptime: 00:08:02
  Accounting status: on/time
  Service accounting session ID: 84:94-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service
  Service State: SvcInactive
  Session ID: 93
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:93-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service6
  Service State: SvcInactive
  Session ID: 95
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:95-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600

```

You can display summary information for all users with a particular access profile. In the following example, only a single user, `vjshah29@example.com`, has the specified access profile, `ce-ppp-profile`:

```

user@host> show accounting pending-accounting-stops ce-ppp-profile

```

Type:	Username:	Session ID:	Service ID:	Service
pppoe	vjshah29@example.com	84		
pppoe	vjshah29@example.com	84	94	cos-service
pppoe	vjshah29@example.com	84	93	filter-service
pppoe	vjshah29@example.com	84	95	filter-service6

You can also display summary information for all subscribers that have accounting-stop messages pending, regardless of access profile. The next example displays information for two users. Because the subscriber larry@example.com is not shown in the previous example, he must have a different access profile than vjshah29@example.com, even though he has received the same services.

```
user@host> show accounting pending-accounting-stops terse
```

Type:	Username:	Session ID:	Service ID:	Service
pppoe	vjshah29@example.com	84		
pppoe	vjshah29@example.com	84	94	cos-service
pppoe	vjshah29@example.com	84	93	filter-service
pppoe	vjshah29@example.com	84	95	filter-service6
pppoe	larry@example.com	85		
pppoe	larry@example.com	85	94	cos-service
pppoe	larry@example.com	85	93	filter-service
pppoe	larry@example.com	85	95	filter-service6

Suspending RADIUS Accounting and Baselineing Accounting Statistics Overview

IN THIS SECTION

- [Sequence of Events During the Suspension, Baselineing, and Resumption of Accounting | 768](#)
- [Guidelines for Accounting Suspension and Baselineing of Statistics | 769](#)
- [Sample Scenarios of Subscriber Accounting Suspension and Baselineing | 769](#)

In certain enterprise provider deployments, maintaining and preserving accounting records might be necessary during a control plane upgrade of a RADIUS accounting server, during an upgrade of the billing system for subscribers, or when RADIUS servers are brought down for maintenance. RADIUS accounting subscriber and service accounting are typically used in such customer topologies for volume-based usage of subscriber traffic and computation of costs. Subscribers might also be billed based on the service level and usage, rather than being charged a set rate regardless of usage.

Starting in Junos OS Release 15.1R4, you can temporarily suspend system-wide accounting until you manually resume accounting. During the suspension period, current subscribers remain logged in, but the subscribers can log out and new subscriber sessions can be initiated. RADIUS Acct-Start, Interim-Update, and Acct-Stop accounting request messages are not generated while accounting is suspended; the router does not send any accounting messages to the RADIUS server. For example, if a subscriber logs out during the suspension, no Acct-Stop request is sent to the server.

After accounting is suspended, all accounting requests are dropped, even if the router is configured to hold the pending accounting messages for up to 24 hours. When accounting resumes, new accounting requests might go into the pending queue, but the requests pending when accounting stopped are no longer available.



NOTE: We do not recommend that operators suspend accounting as a standard practice for system upgrades. However, some operators might find it useful in service provider environments when an upgrade of the server infrastructure is critical and needed immediately.

While accounting is suspended, statistics counters continue to update. You can optionally request a baseline operation to be performed for subscriber and service session time and volume counters. In this case, when accounting is resumed, statistics are reported relative to the baseline values. You can begin the baselining operation only after the suspension starts and before the upgrade begins. You can successfully issue the baseline request only once per suspension. The CLI reports an error if you issue the command again.



NOTE: Statistics are baselined only for subscribers that have interim accounting enabled.

The following RADIUS attributes might be affected for subscribers who are logged in when the baseline is requested and are still logged in when accounting resumes:

- Acct-Session-Time
- Acct-Input-Octets
- Acct-Output-Octets
- Acct-Input-Packets
- Acct-Output-Packets
- Acct-Input-Gigawords
- Acct-Output-Gigawords
- IPv6-Acct-Input-Octets
- IPv6-Acct-Output-Octets
- IPv6-Acct-Input-Packets
- IPv6-Acct-Output-Packets
- IPv6-Acct-Input-Gigawords

- IPv6-Acct-Output-Gigawords

Sequence of Events During the Suspension, Baselineing, and Resumption of Accounting

The following sequence of events occur when you suspend accounting, generate a baseline, and restart accounting processes:

1. Issue the request `network-access aaa accounting suspend` command to suspend accounting.
 - a. A system logging message is generated to indicate that accounting has been suspended.
 - b. All accounting, including accounting-backup-options, is suspended for all accounting servers in all routing contexts.
2. Issue the request `network-access aaa accounting baseline` command to generate a baseline.
 - a. A system logging message is generated to indicate that baselining has started for accounting statistics.
 - b. Time and volume statistics for each subscriber are set to the baseline value. The amount of time that is taken to complete the baseline process is indeterminate, depending on the number of statistical details.
 - c. A system logging message is generated to indicate that baselining has completed.
3. Issue the request `network-access aaa accounting resume` command when baselining is complete to restart accounting processes.
 - a. A system logging message is generated to indicate that accounting has resumed.
 - b. All previously configured accounting options are reenabled.

The baseline operation attempts to baseline the time and volume counters for each subscriber. Subscriber counters are set to baseline values only if interim accounting is enabled for the subscriber by using the `set update-interval minutes` statement at the `[edit access profile profile-name accounting]` hierarchy level. If interim accounting is not enabled for a subscriber, the counters of that corresponding subscriber are not mapped to baseline values.

After the baseline request is executed, an unspecified period of time elapses to baseline all subscriber records. During this interval, statistics for one subscriber can accumulate when the statistical information of another subscriber is being baselined. Sometimes, after baselining starts, counters for some services might be inaccurate and inconsistent due to traffic delivered to a subscriber while the counters of that subscriber are baselined. When the baseline command has been executed, accounting cannot be resumed until the baseline is complete. If you issue the command while accounting is not suspended or while baselining is in progress, the command fails. The command reports an error if the Accounting License is not installed.

Guidelines for Accounting Suspension and Baselineing of Statistics

Keep the following points in mind when you suspend accounting and specify a baseline for statistics:

- Accounting suspension in an environment where thresholds (or quotas) are applicable is not supported. This includes environments where Gx-Plus and Juniper Networks Session and Resource Control (SRC) thresholds or RADIUS session volume quotas are effective for any subscriber. The accounting suspend request fails if any subscriber has thresholds or quotas.
- Activation for threshold (or quota) services is not allowed while accounting is suspended.
- Accounting baselineing is not supported when accounting is not suspended.
- You cannot specify more than one baseline request during an accounting suspension.
- Baselineing for subscribers that are not configured with interim accounting is not supported.
- The time it takes for the baseline operation to complete is indeterminate. It depends on the amount and depth of statistics being collected and is proportional to the number of subscriber and service sessions that are active at the time the baseline is started. The command fails if you attempt to resume accounting while baselineing is still in progress.
- You cannot use the commands to suspend, baseline, or resume accounting during a unified ISSU process. If you attempt to perform a unified ISSU while the baseline is in process, when the chassis daemon state changes to the DAEMON_ISSU_PREPARE state, the authentication and Packet Forwarding Engine processes suspend baselineing on a session boundary and resume after the Routing Engine switchover to the release to which the device is upgraded.
- If a graceful Routing Engine switchover (GRES) occurs while accounting is suspended or baselineing is in progress, the state of suspension or baselineing is preserved after the restart of the router. In such a scenario, accounting is suspended after the reboot of the router and the subscribers for which counters are remaining to be baselined are baselined after the router is online.

Sample Scenarios of Subscriber Accounting Suspension and Baselineing

Consider the following scenario:

1. Interim accounting is configured for subscriber X. It is not configured for subscribers Y and Z.
2. The last interim accounting request sent before accounting is suspended includes statistics for subscriber X; 50,000 octets of traffic have so far been sent for this subscriber. Although 20,000 octets have been sent for subscriber Y and 10,000 octets for subscriber Z, that information has not yet been reported because they do not have interim accounting configured.
3. Accounting is suspended.

4. Baseline begins. The current count for subscriber X is 50,000 octets; this becomes the baseline value for the subscriber. No baseline value is established for subscribers X and Y, because they do not have interim accounting configured.
5. While baselining is in progress, traffic continues to be sent for the three subscribers: 150,000 octets for subscriber X, 80,000 octets for subscriber Y, and 20,000 octets for subscriber Z.
6. Subscriber Z logs out. No Acct-Stop request is sent because accounting is suspended. Consequently, the final accounting statistics are lost for this subscriber.
7. Baseline completes.
8. Accounting resumes.
9. Subscriber X logs out. Although 200,000 total octets were sent for subscriber X, the Acct-Stop record reports only 150,000 octets: 200,000 total octets minus the 50,000 octet baseline.
10. Subscriber Y logs out. Because 100,000 total octets were sent for subscriber Y and there is no baseline value, the Acct-Stop record reports the total of 100,000 octets.

Table 19 on page 770 summarizes this scenario.

Table 19: Summary of Accounting Suspension and Baselining Scenario

Subscriber	Interim Accounting configured	Octets Before Suspension	Octets After Baselining Starts	Total Octets	Octets in Acct-Stop When Accounting Resumes
X	Yes	50,000	150,000	200,000	150,000
Y	No	20,000	80,000	100,000	100,000
Z	No	10,000	20,000	30,000	n/a

Configuring RADIUS Accounting Suspension and Baselining Accounting Statistics

You can temporarily suspend system-wide accounting for the duration of a system upgrade or maintenance action, until you manually resume accounting. During the suspension period, current subscribers remain logged in, but the subscribers can log out and new subscriber sessions can be initiated. RADIUS Acct-Start, Interim-Update, and Acct-Stop messages are not generated while accounting is suspended; the router does not send any accounting messages to the RADIUS server. For example, if a subscriber logs out during the suspension, no Acct-Stop is sent to the server.



NOTE: We do not recommend that operators suspend accounting as a standard practice for system upgrades. However, some operators might find it useful in service provider environments when an upgrade of the server infrastructure is critical and needed immediately.

To configure the suspension of accounting processes, create a baseline of the statistics after accounting is halted, and resume accounting after the baselining process is completed:

1. Suspend subscriber accounting.

```
user@host> request network-access aaa accounting suspend
```

A syslog message is generated to indicate that accounting is suspended. All accounting (including accounting-backup-options) is suspended for all accounting servers and all routing contexts.

2. (Optional) Begin baselining accounting statistics for subscribers that have interim accounting configured.

```
user@host> request network-access aaa accounting baseline
```

The router implements the baseline by reading and storing the statistics when the baseline is set. The baseline values are subtracted when you retrieve baseline-relative statistics after accounting resumes. A syslog message is generated to indicate the start of baselining. Time and volume statistics for each subscriber are set to the baseline value. The amount of time that is taken to complete the baseline process might vary, depending on the number of statistical details. A syslog message is generated when the baselining of statistics completes.

3. Resume accounting after baselining completes.

```
user@host> request network-access aaa accounting resume
```

A syslog message is generated to indicate that accounting has resumed. All the previously configured accounting options are reenabled.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.1R4	Starting in Junos OS Release 15.1R4, you can temporarily suspend system-wide accounting until you manually resume accounting.

RELATED DOCUMENTATION

AAA Service Framework Overview 2
RADIUS Authentication and Accounting Basic Configuration 721
Standard and Vendor-Specific RADIUS Attributes 3

Verifying and Managing Subscriber AAA Information

IN THIS SECTION

- [Purpose | 772](#)
- [Action | 772](#)

Purpose

View or clear subscriber access statistics and information.

Action

- To display subscriber AAA statistics:

```
user@host> show network-access aaa statistics
```

```
user@host> show network-access aaa statistics authentication
```

- To display RADIUS server status and information:

```
user@host> show network-access aaa radius-servers
```

- To display subscriber access AAA information:

```
user@host> show network-access aaa subscribers
```

- To display subscriber session information:

```
user@host> show network-access aaa subscribers session-id session-id
```

- To clear subscriber access statistics and to log out specific subscribers:

```
user@host> clear network-access aaa subscriber
```

You can specify the subscriber with the `username username` option or the `session-id identifier` option. In either case, specify `reconnect` to attempt to reconnect the subscriber session after it is completely logged out.

- To clear blocked request statistics to debug session limits for all usernames across all access profiles:

```
user@host> clear network-access aaa subscriber session-limit-per-username
```

- To clear blocked request statistics to debug session limits for a specific username across all access profiles:

```
user@host> clear network-access aaa subscriber session-limit-per-username username username
```

- To clear blocked request statistics to debug session limits for all usernames in a specific access profile:

```
user@host> clear network-access aaa subscriber session-limit-per-username access-profile profile-name
```

- To clear blocked request statistics to debug session limits for a specific username in a specific access profile:

```
user@host> clear network-access aaa subscriber session-limit-per-username username username  
access-profile profile-name
```

- To clear AAA accounting statistics:

```
user@host> clear network-access aaa statistics accounting
```

- To clear AAA address-assignment statistics for a client:

```
user@host> clear network-access aaa statistics address-assignment client
```

- To clear AAA address-assignment pool statistics:

```
user@host> clear network-access aaa statistics address-assignment pool pool-name
```

- To clear AAA authentication statistics:

```
user@host> clear network-access aaa statistics authentication
```

RELATED DOCUMENTATION

| [RADIUS Servers and Parameters for Subscriber Access](#) | 646

Session Termination Causes and RADIUS Termination Cause Codes

IN THIS SECTION

- [Understanding Session Termination Causes and RADIUS Termination Cause Codes](#) | 775
- [Mapping Session Termination Causes to Custom Termination Cause Codes](#) | 777

Understanding Session Termination Causes and RADIUS Termination Cause Codes

IN THIS SECTION


- [Benefits of Session and Service Termination Cause Codes | 777](#)

When a RADIUS Acct-Stop message is issued as a result of the termination of a subscriber session or service session, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. Default mappings exist for AAA, DHCP, L2TP, PPP, and VLAN subscriber and service sessions. When a subscriber or service session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute.

You can use the logged information to help monitor and troubleshoot the events. For example, the AAA termination causes include session and service terminations as well as access denials. You might want to route the access failures to a team that monitors attempts to hack the network, the timeout failures to a AAA server team, and resource failures to a team that manages the routers.

Because there are many different Junos OS internal identifiers for termination causes and only 18 standard code values defined in the RFC, by default a given code value can map to multiple identifiers. Instead of using the default code values, you can optionally map any of the internally defined termination causes to any 32-bit number (1 through 4,294,967,295). The flexibility of customized mapping greatly increases the possibilities for fine-grained analytics and failure tracking.



NOTE: A single mapping for RADIUS account termination is shared by all clients.

[Table 20 on page 775](#) lists the RFC-defined standard RADIUS Acct-Terminate-Cause codes and the corresponding causes.

Table 20: RFC-Defined Code Values and Termination Causes

Code Value	Termination Cause	Description
1	User Request	User initiated the disconnect (logout).

Table 20: RFC-Defined Code Values and Termination Causes (Continued)

Code Value	Termination Cause	Description
2	Lost Carrier	DCD was dropped on the port.
3	Lost Service	Service can no longer be provided; for example, the user's connection to a host was interrupted.
4	Idle Timeout	Idle timer expired.
5	Session Timeout	Subscriber reached the maximum continuous time allowed for the service or session.
6	Admin Reset	System administrator reset the port or session.
7	Admin Reboot	System administrator terminated the session on the NAS; for example, prior to rebooting the NAS.
8	Port Error	NAS detected an error on the port that required ending the session.
9	NAS Error	NAS detected an error (other than on the port) that required ending the session.
10	NAS Request	NAS ended the session for a non-error reason.
11	NAS Reboot	NAS ended the session due to a non-administrative reboot.
12	Port Unneeded	NAS ended the session because the resource usage fell below the low threshold; for example, the bandwidth-on-demand algorithm determined that the port was no longer needed.
13	Port Preempted	NAS ended the session to allocate the port to a higher-priority use.
14	Port Suspended	NAS ended the session to suspend a virtual session.

Table 20: RFC-Defined Code Values and Termination Causes (Continued)

Code Value	Termination Cause	Description
15	Service Unavailable	NAS was unable to provide the requested service.
16	Callback	NAS is terminating the current session in order to perform callback for a new session.
17	User Error	Error in the user input caused the session to be terminated.
18	Host Request	Login host terminated the session normally.

Benefits of Session and Service Termination Cause Codes

- Termination cause codes mapped to Junos OS internal identifiers can help you monitor, analyze, and troubleshoot the events that resulted in termination of subscriber sessions or service sessions.
- Customized mappings enable you to map internal termination cause identifiers for termination cause codes to a code value of your choosing for more fine-grained tracking and analysis of termination events.

Mapping Session Termination Causes to Custom Termination Cause Codes

By default, Junos OS uses the RFC-defined termination cause codes for the internal identifiers that identify the causes of session termination and that are reported in the RADIUS Acct-Terminate-Cause attribute (49). Internal identifiers are available for AAA, DHCP, L2TP, PPP, and VLAN subscriber and service session failures. When a subscriber or service session is terminated or denied, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. The Acct-Terminate-Cause attribute is included in RADIUS Acct-Stop messages. You can use the logged information to help monitor and troubleshoot terminated sessions.

You can optionally create customized mappings between any of the internal termination cause identifiers for the protocol and termination cause codes. You can specify any 32-bit value for the code, enabling you to track and analyze particular termination events at a more fine-grained level.

To configure customized mappings between a termination cause and a RADIUS cause code:

1. Edit the access hierarchy.

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

2. Edit the terminate-code statement.



NOTE: Termination cause codes do not appear as options on platforms where they are not supported.

```
[edit access]
user@host# edit terminate-code
```

3. Specify the protocol option (aaa (deny | service-shutdown | shutdown) | dhcp | l2tp | ppp | vlan) that you want to modify.

```
[edit access terminate-code]
user@host# edit protocol-option
```

4. Specify an existing termination cause that you want to remap.

```
[edit access terminate-code protocol-option]
user@host# edit term-reason
```



NOTE: Attempts to remap a termination cause to its default code value are rejected by the CLI. You must delete a custom mapping to restore the default mapping.

5. Specify the RADIUS termination cause code value (from 1 through 4,294,967,295) that you want to map to the termination cause.

```
[edit access terminate-code protocol-option term-reason]
user@host# set radius term-cause
```

Use the `show network-access aaa terminate-code` command to display the mapping between AAA termination causes and cause code values.

RELATED DOCUMENTATION

[AAA Termination Causes and Code Values | 779](#)

[DHCP Termination Causes and Code Values | 781](#)

[L2TP Termination Causes and Code Values | 783](#)

[PPP Termination Causes and Code Values | 809](#)

[VLAN Termination Causes and Code Values | 822](#)

AAA Termination Causes and Code Values

When a AAA event terminates a subscriber or service session, causing a RADIUS Acct-Stop message to be issued, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. When a subscriber or service session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot terminated sessions.

[Table 21 on page 779](#) lists the default mapping between the internal identifier for AAA termination causes and the code values that represent them in the RADIUS Acct-Terminate-Cause attribute (49).



NOTE: You can remap the internal identifiers to a custom code value in the range 1 through 4,294,967,295 by using the `terminate-code` statement at the [edit access] hierarchy level. You can view the current mapping by issuing the `show network-access terminate-code aaa detail` command.

Table 21: Default Mapping Between AAA Termination Causes and Code Values

Internal AAA Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	Description
deny-authentication-denied	17	Subscriber access denied due to authentication failure.

Table 21: Default Mapping Between AAA Termination Causes and Code Values (Continued)

Internal AAA Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	Description
deny-no-resources	10	Subscriber access denied for reasons such as no RADIUS server exists.
deny-server-request-timeout	17	Subscriber access denied because the BNG retried the Access-Request to the authentication server for the configured number of retries without receiving a response.
service-shutdown-network-logout	6	Service session termination initiated by deactivation of a family (network), typically triggered by termination of the corresponding Layer 3 access protocol.
service-shutdown-remote-reset	10	Service session termination initiated by an external authority, such as a CoA service deactivation.
service-shutdown-subscriber-logout	Inherited from the parent subscriber session.	Overrides the default value. This code is displayed only when you map it to a custom value.
service-shutdown-time-limit	5	Service session termination initiated because the service time limit was reached.
service-shutdown-volume-limit	10	Service session termination initiated because the service traffic volume limit was reached.
shutdown-administrative-reset	6	Session has been terminated by a local CLI command (such as the <code>clear dhcp server binding</code> command) or to clean up dynamic VLAN configured with "remove-when-no-subscribers" when there is no successful subscriber connections over that VLAN within 30 seconds after its creation.)

Table 21: Default Mapping Between AAA Termination Causes and Code Values *(Continued)*

Internal AAA Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	Description
shutdown-idle-timeout	4	Session has been idle for a period equal to or longer than the configured timeout value. This value is set with the CLI or by RADIUS attribute.
shutdown-reassign-on-match	10	Session is terminated to allow a second session to replace the terminated session. This occurs only when both sessions are allocated the same static IP address by means of the RADIUS Framed-IP-Address attribute (8). This behavior enables a customer to reconnect with a new session after dropping off the original session, even though the original session is still up.
shutdown-remote-reset	10	Session has been terminated by a remote service, such as a RADIUS Disconnect-Request or Diameter Abort-Session-Request messages.
shutdown-session-timeout	5	Session has been active for a period equal to or longer than the configured timeout value. This value is set with the CLI or by RADIUS attribute.

RELATED DOCUMENTATION


[Session Termination Causes and RADIUS Termination Cause Codes](#) | 774

DHCP Termination Causes and Code Values

When a DHCP event terminates a subscriber session, causing a RADIUS Acct-Stop message to be issued, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. When a subscriber session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot terminated sessions.

Table 22 on page 782 lists the default mapping between the internal identifier for DHCP termination causes and the code values that represent them in the RADIUS Acct-Terminate-Cause attribute (49).



NOTE: You can remap the internal identifiers to a custom code value in the range 1 through 4,294,967,295 by using the `terminate-code` statement at the `[edit access]` hierarchy level. You can view the current mapping by issuing the `show network-access terminate-code dhcp detail` command.

Table 22: Default Mapping Between DHCP Termination Causes and Code Values

Internal DHCP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
client request	1	User Request
lost-carrier	2	Lost Carrier
nak	15	Service Unavailable
nas logout	10	NAS Request
no offers	4	Idle Timeout

RELATED DOCUMENTATION

| [Session Termination Causes and RADIUS Termination Cause Codes](#) | 774

L2TP Termination Causes and Code Values

When an L2TP event terminates a subscriber session, causing a RADIUS Acct-Stop message to be issued, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. When a subscriber session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot terminated sessions.

[Table 23 on page 783](#) lists the default mapping between the internal identifier for L2TP termination causes and the code values that represent them in the RADIUS Acct-Terminate-Cause attribute (49).



NOTE: You can remap the internal identifiers to a custom code value in the range 1 through 4,294,967,295 by using the `terminate-code` statement at the `[edit access]` hierarchy level. You can view the current mapping by issuing the `show network-access terminate-code l2tp detail` command.

Table 23: Default Mapping Between L2TP Termination Causes and Code Values

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
issu in progress	9	NAS Error
session access interface down	8	Port Error
session admin close	6	Admin Reset
session admin drain	6	Admin Reset
session call down	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session call failed	15	Service Unavailable
session create failed limit reached	9	NAS Error
session create failed no resources	9	NAS Error
session create failed single shot tunnel already fired	9	NAS Error
session create failed too busy	9	NAS Error
session failover protocol resync disconnect	6	Admin Reset
session hardware unavailable	8	Port Error
session no resources server port	9	NAS Error
session not ready	9	NAS Error
session rx cdn	10	NAS Request
session rx cdn avp bad hidden	10	NAS Request
session rx cdn avp bad value assigned session id	10	NAS Request
session rx cdn avp duplicate value assigned session id	10	NAS Request
session rx cdn avp malformed bad length	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx cdn avp malformed truncated	10	NAS Request
session rx cdn avp missing mandatory assigned session id	10	NAS Request
session rx cdn avp missing mandatory result code	10	NAS Request
session rx cdn avp missing random vector	10	NAS Request
session rx cdn avp missing secret	10	NAS Request
session rx cdn avp unknown	10	NAS Request
session rx cdn no resources	10	NAS Request
session rx iccn avp bad hidden	10	NAS Request
session rx iccn avp bad value framing type	10	NAS Request
session rx iccn avp bad value proxy authen type	10	NAS Request
session rx iccn avp bad value unsupported proxy authen type	10	NAS Request
session rx iccn avp malformed bad length	10	NAS Request
session rx iccn avp malformed truncated	10	NAS Request
session rx iccn avp missing mandatory connect speed	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx iccn avp missing mandatory framing type	10	NAS Request
session rx iccn avp missing mandatory proxy authen challenge	10	NAS Request
session rx iccn avp missing mandatory proxy authen id	10	NAS Request
session rx iccn avp missing mandatory proxy authen name	10	NAS Request
session rx iccn avp missing mandatory proxy authen response	10	NAS Request
session rx iccn avp missing random vector	10	NAS Request
session rx iccn avp missing secret	10	NAS Request
session rx iccn avp unknown	10	NAS Request
session rx iccn no resources	10	NAS Request
session rx iccn unexpected	10	NAS Request
session rx icrp avp bad hidden	10	NAS Request
session rx icrp avp bad value assigned session id	10	NAS Request
session rx icrp avp duplicate value assigned session id	10	NAS Request
session rx icrp avp malformed bad length	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx icrp avp malformed truncated	10	NAS Request
session rx icrp avp missing mandatory assigned session id	10	NAS Request
session rx icrp avp missing random vector	10	NAS Request
session rx icrp avp missing secret	10	NAS Request
session rx icrp avp unknown	10	NAS Request
session rx icrp no resources	10	NAS Request
session rx icrp unexpected	10	NAS Request
session rx icrq admin close	6	Admin Reset
session rx icrq authenticate failed host	10	NAS Request
session rx icrq avp bad hidden	10	NAS Request
session rx icrq avp bad value assigned session id	10	NAS Request
session rx icrq avp bad value bearer type	10	NAS Request
session rx icrq avp bad value cisco nas port	10	NAS Request
session rx icrq avp duplicate value assigned session id	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx icrq avp malformed bad length	10	NAS Request
session rx icrq avp malformed truncated	10	NAS Request
session rx icrq avp missing mandatory assigned session id	10	NAS Request
session rx icrq avp missing mandatory call serial number	10	NAS Request
session rx icrq avp missing random vector	10	NAS Request
session rx icrq avp missing secret	10	NAS Request
session rx icrq avp unknown	10	NAS Request
session rx icrq no resources	10	NAS Request
session rx icrq unexpected	10	NAS Request
session rx occn avp bad hidden	10	NAS Request
session rx occn avp bad value framing type	10	NAS Request
session rx occn avp malformed bad length	10	NAS Request
session rx occn avp malformed truncated	10	NAS Request
session rx occn avp missing mandatory connect speed	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx occn avp missing mandatory framing type	10	NAS Request
session rx occn avp missing random vector	10	NAS Request
session rx occn avp missing secret	10	NAS Request
session rx occn avp unknown	10	NAS Request
session rx occn no resources	10	NAS Request
session rx occn unexpected	10	NAS Request
session rx ocrp avp bad hidden	10	NAS Request
session rx ocrp avp bad value assigned session id	10	NAS Request
session rx ocrp avp duplicate value assigned session id	10	NAS Request
session rx ocrp avp malformed bad length	10	NAS Request
session rx ocrp avp malformed truncated	10	NAS Request
session rx ocrp avp missing mandatory assigned session id	10	NAS Request
session rx ocrp avp missing random vector	10	NAS Request
session rx ocrp avp missing secret	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx ocrp avp unknown	10	NAS Request
session rx ocrp no resources	10	NAS Request
session rx ocrp unexpected	10	NAS Request
session rx ocrq admin close	10	Admin Reset
session rx ocrq authenticate failed host	10	NAS Request
session rx ocrq avp bad hidden	10	NAS Request
session rx ocrq avp bad value assigned session id	10	NAS Request
session rx ocrq avp bad value bearer type	10	NAS Request
session rx ocrq avp bad value framing type	10	NAS Request
session rx ocrq avp duplicate value assigned session id	10	NAS Request
session rx ocrq avp malformed bad length	10	NAS Request
session rx ocrq avp malformed truncated	10	NAS Request
session rx ocrq avp missing mandatory assigned session id	10	NAS Request
session rx ocrq avp missing mandatory bearer type	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx ocrq avp missing mandatory call serial number	10	NAS Request
session rx ocrq avp missing mandatory called number	10	NAS Request
session rx ocrq avp missing mandatory framing type	10	NAS Request
session rx ocrq avp missing mandatory maximum bps	10	NAS Request
session rx ocrq avp missing mandatory minimum bps	10	NAS Request
session rx ocrq avp missing random vector	10	NAS Request
session rx ocrq avp missing secret	10	NAS Request
session rx ocrq avp unknown	10	NAS Request
session rx ocrq no resources	10	NAS Request
session rx ocrq unexpected	10	NAS Request
session rx ocrq unsupported	9	NAS Error
session rx sli avp bad hidden	10	NAS Request
session rx sli avp bad value accm	10	NAS Request
session rx sli avp malformed bad length	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx sli avp malformed truncated	10	NAS Request
session rx sli avp missing mandatory accm	10	NAS Request
session rx sli avp missing random vector	10	NAS Request
session rx sli avp missing secret	10	NAS Request
session rx sli avp unknown	10	NAS Request
session rx sli no resources	10	NAS Request
session rx unexpected packet lac incoming	10	NAS Request
session rx unexpected packet lac outgoing	10	NAS Request
session rx unexpected packet lns incoming	10	NAS Request
session rx unexpected packet lns outgoing	10	NAS Request
session rx unknown session id	10	NAS Request
session rx wen avp bad hidden	10	NAS Request
session rx wen avp malformed bad length	10	NAS Request
session rx wen avp malformed truncated	10	NAS Request

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx wen avp missing mandatory call errors	10	NAS Request
session rx wen avp missing random vector	10	NAS Request
session rx wen avp missing secret	10	NAS Request
session rx wen avp unknown	10	NAS Request
session rx wen no resources	10	NAS Request
session timeout connection	10	NAS Request
session timeout inactivity	4	idle timeout
session timeout session	5	session timeout
session timeout upper create	9	NAS Error
session transmit speed unavailable	9	NAS error
session tunnel down	15	Service Unavailable
session tunnel failed	15	Service Unavailable
session tunnel switch profile deleted	6	Admin Reset
session tunneled interface down	8	Port Error

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session unknown cause	9	NAS Error
session upper create failed	9	NAS Error
session upper removed	15	Service Unavailable
session warmstart not operational	15	Service Unavailable
session warmstart recovery error	15	Service Unavailable
session warmstart upper not restacked	10	NAS request
tunnel admin close	6	Admin Reset
tunnel admin drain	6	Admin Reset
tunnel control channel failed	15	Service Unavailable
tunnel created no sessions	1	User Request
tunnel destination address changed	6	Admin Reset
tunnel destination down	10	NAS Request
tunnel failover protocol no resources for recovery tunnel	15	Service Unavailable
tunnel failover protocol no resources for session resync	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel failover protocol not supported	15	Service Unavailable
tunnel failover protocol not supported by peer	15	Service Unavailable
tunnel failover protocol recovery control channel failed	15	Service Unavailable
tunnel failover protocol recovery tunnel failed	15	Service Unavailable
tunnel failover protocol recovery tunnel finished	1	User Request
tunnel failover protocol recovery tunnel primary down	1	User Request
tunnel failover protocol session resync failed	15	Service Unavailable
tunnel host profile changed	6	Admin Reset
tunnel host profile deleted	6	Admin Reset
tunnel rx sccn authenticate failed challenge	17	User Error
tunnel rx sccn avp bad hidden	15	Service Unavailable
tunnel rx sccn avp bad value challenge response	15	Service Unavailable
tunnel rx sccn avp malformed bad length	15	Service Unavailable
tunnel rx sccn avp malformed truncated	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx scccn avp missing challenge response	17	User Error
tunnel rx scccn avp missing random vector	15	Service Unavailable
tunnel rx scccn avp missing secret	15	Service Unavailable
tunnel rx scccn avp unexpected challenge response	15	Service Unavailable
tunnel rx scccn avp unknown	15	Service Unavailable
tunnel rx scccn no resources	15	Service Unavailable
tunnel rx scccn session id not null	15	Service Unavailable
tunnel rx scccn unexpected	15	Service Unavailable
tunnel rx sccrp authenticate failed challenge	17	User Error
tunnel rx sccrp authenticate failed host	17	User Error
tunnel rx sccrp avp bad hidden	15	Service Unavailable
tunnel rx sccrp avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx sccrp avp bad value bearer capabilities	15	Service Unavailable
tunnel rx sccrp avp bad value challenge	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrp avp bad value challenge response	15	Service Unavailable
tunnel rx sccrp avp bad value failover capability	15	Service Unavailable
tunnel rx sccrp avp bad value framing capabilities	15	Service Unavailable
tunnel rx sccrp avp bad value protocol version	15	Service Unavailable
tunnel rx sccrp avp bad value receive window size	15	Service Unavailable
tunnel rx sccrp avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx sccrp avp malformed bad length	15	Service Unavailable
tunnel rx sccrp avp malformed truncated	15	Service Unavailable
tunnel rx sccrp avp missing challenge response	17	User Error
tunnel rx sccrp avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx sccrp avp missing mandatory framing capabilities	15	Service Unavailable
tunnel rx sccrp avp missing mandatory host name	15	Service Unavailable
tunnel rx sccrp avp missing mandatory protocol version	15	Service Unavailable
tunnel rx sccrp avp missing random vector	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrp avp missing secret	15	Service Unavailable
tunnel rx sccrp avp unexpected challenge response	15	Service Unavailable
tunnel rx sccrp avp unexpected challenge without secret	15	Service Unavailable
tunnel rx sccrp avp unknown	15	Service Unavailable
tunnel rx sccrp no resources	15	Service Unavailable
tunnel rx sccrp session id not null	15	Service Unavailable
tunnel rx sccrp unexpected	15	Service Unavailable
tunnel rx sccrq admin close	6	Admin Reset
tunnel rx sccrq authenticate failed host	17	User Error
tunnel rx sccrq avp bad hidden	15	Service Unavailable
tunnel rx sccrq avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx sccrq avp bad value bearer capabilities	15	Service Unavailable
tunnel rx sccrq avp bad value challenge	15	Service Unavailable
tunnel rx sccrq avp bad value failover capability	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrq avp bad value framing capabilities	15	Service Unavailable
tunnel rx sccrq avp bad value protocol version	15	Service Unavailable
tunnel rx sccrq avp bad value receive window size	15	Service Unavailable
tunnel rx sccrq avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx sccrq avp malformed bad length	15	Service Unavailable
tunnel rx sccrq avp malformed truncated	15	Service Unavailable
tunnel rx sccrq avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx sccrq avp missing mandatory framing capabilities	15	Service Unavailable
tunnel rx sccrq avp missing mandatory host name	15	Service Unavailable
tunnel rx sccrq avp missing mandatory protocol version	15	Service Unavailable
tunnel rx sccrq avp missing random vector	15	Service Unavailable
tunnel rx sccrq avp missing secret	15	Service Unavailable
tunnel rx sccrq avp unexpected challenge without secret	15	Service Unavailable
tunnel rx sccrq avp unknown	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrq bad address	15	Service Unavailable
tunnel rx sccrq no resources	15	Service Unavailable
tunnel rx sccrq no resources max tunnels	15	Service Unavailable
tunnel rx sccrq session id not null	15	Service Unavailable
tunnel rx sccrq unexpected	15	Service Unavailable
tunnel rx stopccn	1	User Request
tunnel rx stopccn avp bad hidden	15	Service Unavailable
tunnel rx stopccn avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx stopccn avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx stopccn avp malformed bad length	15	Service Unavailable
tunnel rx stopccn avp malformed truncated	15	Service Unavailable
tunnel rx stopccn avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx stopccn avp missing mandatory result code	15	Service Unavailable
tunnel rx stopccn avp missing random vector	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx stopccn avp missing secret	15	Service Unavailable
tunnel rx stopccn avp unknown	15	Service Unavailable
tunnel rx stopccn no resources	15	Service Unavailable
tunnel rx stopccn session id not null	15	Service Unavailable
tunnel rx frs avp malformed truncated	15	Service Unavailable
tunnel rx frs avp missing mandatory failover session state	15	Service Unavailable
tunnel rx frs avp missing random vector	15	Service Unavailable
tunnel rx frs avp missing secret	15	Service Unavailable
tunnel rx frs avp unknown	15	Service Unavailable
tunnel rx frs no resources	15	Service Unavailable
tunnel rx frs session id not null	15	Service Unavailable
tunnel rx fsq avp bad hidden	15	Service Unavailable
tunnel rx fsq avp malformed bad length	15	Service Unavailable
tunnel rx fsq avp malformed truncated	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx fsq avp missing mandatory failover session state	15	Service Unavailable
tunnel rx fsq avp missing random vector	15	Service Unavailable
tunnel rx fsq avp missing secret	15	Service Unavailable
tunnel rx fsq avp unknown	15	Service Unavailable
tunnel rx fsq no resources	15	Service Unavailable
tunnel rx fsq session id not null	15	Service Unavailable
tunnel rx fsr avp bad hidden	15	Service Unavailable
tunnel rx fsr avp malformed bad length	15	Service Unavailable
tunnel rx unexpected packet	15	Service Unavailable
tunnel rx unexpected packet for session	15	Service Unavailable
tunnel rx unknown packet message type indecipherable	15	Service Unavailable
tunnel rx unknown packet message type unrecognized	15	Service Unavailable
tunnel rx recovery sccn authenticate failed challenge	17	User Error
tunnel rx recovery sccn avp bad hidden	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery scccn avp bad value challenge response	15	Service Unavailable
tunnel rx recovery scccn avp malformed bad length	15	Service Unavailable
tunnel rx recovery scccn avp malformed truncated	15	Service Unavailable
tunnel rx recovery scccn avp missing challenge response	17	User Error
tunnel rx recovery scccn avp missing random vector	15	Service Unavailable
tunnel rx recovery scccn avp missing secret	15	Service Unavailable
tunnel rx recovery scccn avp unexpected challenge response	15	Service Unavailable
tunnel rx recovery scccn avp unknown	15	Service Unavailable
tunnel rx recovery scccn no resources	15	Service Unavailable
tunnel rx recovery scccn session id not null	15	Service Unavailable
tunnel rx recovery sccrp authenticate failed challenge	17	User Error
tunnel rx recovery sccrp avp bad hidden	15	Service Unavailable
tunnel rx recovery sccrp avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrp avp bad value bearer capabilities	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery sccrp avp bad value challenge	15	Service Unavailable
tunnel rx recovery sccrp avp bad value challenge response	15	Service Unavailable
tunnel rx recovery sccrp avp bad value framing capabilities	15	Service Unavailable
tunnel rx recovery sccrp avp bad value protocol version	15	Service Unavailable
tunnel rx recovery sccrp avp bad value receive window size	15	Service Unavailable
tunnel rx recovery sccrp avp bad value suggested control sequence	15	Service Unavailable
tunnel rx recovery sccrp avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrp avp malformed bad length	15	Service Unavailable
tunnel rx recovery sccrp avp malformed truncated	15	Service Unavailable
tunnel rx recovery sccrp avp mismatched host name	15	Service Unavailable
tunnel rx recovery sccrp avp mismatched vendor name	15	Service Unavailable
tunnel rx recovery sccrp avp missing challenge response	17	User Error
tunnel rx recovery sccrp avp missing mandatory assigned tunnel id	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery sccrp avp missing mandatory framing capabilities	15	Service Unavailable
tunnel rx recovery sccrp avp missing mandatory host name	15	Service Unavailable
tunnel rx recovery sccrp avp missing mandatory protocol version	15	Service Unavailable
tunnel rx recovery sccrp avp missing random vector	15	Service Unavailable
tunnel rx recovery sccrp avp missing secret	15	Service Unavailable
tunnel rx recovery sccrp avp unexpected challenge response	15	Service Unavailable
tunnel rx recovery sccrp avp unexpected challenge without secret	15	Service Unavailable
tunnel rx recovery sccrp avp unknown	15	Service Unavailable
tunnel rx recovery sccrp no resources	15	Service Unavailable
tunnel rx recovery sccrp session id not null	15	Service Unavailable
tunnel rx recovery sccrq admin close	6	Admin Reset
tunnel rx recovery sccrq avp bad hidden	15	Service Unavailable
tunnel rx recovery sccrq avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrq avp bad value bearer capabilities	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery sccrq avp bad value challenge	15	Service Unavailable
tunnel rx recovery sccrq avp bad value framing capabilities	15	Service Unavailable
tunnel rx recovery sccrq avp bad value protocol version	15	Service Unavailable
tunnel rx recovery sccrq avp bad value receive window size	15	Service Unavailable
tunnel rx recovery sccrq avp bad value tunnel recovery	15	Service Unavailable
tunnel rx recovery sccrq avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrq avp duplicate value tie breaker	15	Service Unavailable
tunnel rx recovery sccrq avp malformed bad length	15	Service Unavailable
tunnel rx recovery sccrq avp malformed truncated	15	Service Unavailable
tunnel rx recovery sccrq avp mismatched host name	15	Service Unavailable
tunnel rx recovery sccrq avp mismatched vendor name	15	Service Unavailable
tunnel rx recovery sccrq avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrq avp missing mandatory framing capabilities	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery sccrq avp missing mandatory host name	15	Service Unavailable
tunnel rx recovery sccrq avp missing mandatory protocol version	15	Service Unavailable
tunnel rx recovery sccrq avp missing mandatory tunnel recovery	15	Service Unavailable
tunnel rx recovery sccrq avp missing random vector	15	Service Unavailable
tunnel rx recovery sccrq avp missing secret	15	Service Unavailable
tunnel rx recovery sccrq avp missing tie breaker	15	Service Unavailable
tunnel rx recovery sccrq avp unexpected challenge without secret	15	Service Unavailable
tunnel rx recovery sccrq avp unknown	15	Service Unavailable
tunnel rx recovery sccrq no resources	15	Service Unavailable
tunnel rx recovery sccrq session id not null	15	Service Unavailable
tunnel rx recovery sccrq tunnel id not null	15	Service Unavailable
tunnel rx recovery stopccn avp bad hidden	15	Service Unavailable
tunnel rx recovery stopccn avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx recovery stopccn avp duplicate value assigned tunnel id	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values *(Continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery stopccn avp malformed bad length	15	Service Unavailable
tunnel rx recovery stopccn avp malformed truncated	15	Service Unavailable
tunnel rx recovery stopccn avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx recovery stopccn avp missing mandatory result code	15	Service Unavailable
tunnel rx recovery stopccn avp missing random vector	15	Service Unavailable
tunnel rx recovery stopccn avp missing secret	15	Service Unavailable
tunnel rx recovery stopccn avp unknown	15	Service Unavailable
tunnel rx recovery stopccn no resources	15	Service Unavailable
tunnel rx recovery stopccn session id not null	15	Service Unavailable
tunnel rx recovery unexpected packet	15	Service Unavailable
tunnel rx recovery unknown packet message type indecipherable	15	Service Unavailable
tunnel rx recovery unknown packet message type unrecognized	15	Service Unavailable
tunnel rx session packet null sid invalid	15	Service Unavailable
tunnel rx session packet null sid without assigned session id	15	Service Unavailable

Table 23: Default Mapping Between L2TP Termination Causes and Code Values (*Continued*)

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel timeout connection	15	Service Unavailable
tunnel timeout connection recovery tunnel	15	Service Unavailable
tunnel timeout idle	1	User Request
tunnel unknown cause	9	NAS Error
tunnel warmstart not operational	15	Service Unavailable
tunnel warmstart recovery error	15	Service Unavailable

RELATED DOCUMENTATION

[Session Termination Causes and RADIUS Termination Cause Codes](#) | 774

PPP Termination Causes and Code Values

When a PPP event terminates a subscriber session, causing a RADIUS Acct-Stop message to be issued, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, defines the standard mapping between 18 termination causes and code values.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. When a subscriber session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot terminated sessions.

Table 24 on page 810 lists the default mapping between the internal identifier for PPP termination causes and the code values that represent them in the RADIUS Acct-Terminate-Cause attribute (49).



NOTE: You can remap the internal identifiers to a custom code value in the range 1 through 4,294,967,295 by using the `terminate-code` statement at the [edit access] hierarchy level. You can view the current mapping by issuing the `show network-access terminate-code ppp detail` command.

Table 24 on page 810 lists the default PPP terminate mappings. The table indicates the supported PPP terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 24: Default Mapping Between PPP Termination Causes and Code Values

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
admin logout	10	NAS Request
authenticate authenticator timeout	17	User Error
authenticate challenge timeout	10	NAS Request
authenticate chap no resources	10	NAS Request
authenticate chap peer authenticator timeout	17	User Error
authenticate deny by peer	17	User Error
authenticate inactivity timeout	4	Idle Timeout
authenticate max requests	10	NAS Request
authenticate no authenticator	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
authenticate pap peer authenticator timeout	17	User Error
authenticate pap request timeout	10	NAS Request
authenticate Session Timeout	5	Session Timeout
authenticate too many requests	10	NAS Request
authenticate tunnel fail immediate	10	NAS Request
authenticate tunnel unsupported tunnel type	10	NAS Request
bundle fail create	10	NAS Request
bundle fail engine add	10	NAS Request
bundle fail fragment size mismatch	10	NAS Request
bundle fail fragmentation location	10	NAS Request
bundle fail fragmentation mismatch	10	NAS Request
bundle fail join	10	NAS Request
bundle fail link selection mismatch	10	NAS Request
bundle fail local mpd not set yet	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
bundle fail local mrru mismatch	10	NAS Request
bundle fail local mru mismatch	10	NAS Request
bundle fail peer mrru mismatch	10	NAS Request
bundle fail reassembly location	10	NAS Request
bundle fail reassembly mismatch	10	NAS Request
bundle fail record network	10	NAS Request
bundle fail server location mismatch	10	NAS Request
bundle fail static link	10	NAS Request
failover during authentication	6	Admin Reset
interface admin disable	6	Admin Reset
interface down	2	Lost Carrier
interface no hardware	8	Port Error
ip admin disable	10	NAS Request
ip inhibited by authentication	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values *(Continued)*

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
ip link down	10	NAS Request
ip max configure exceeded	10	NAS Request
ip no local ip address	10	NAS Request
ip no local ip address mask	10	NAS Request
ip no local primary dns address	10	NAS Request
ip no local primary nbns address	10	NAS Request
ip no local secondary dns address	10	NAS Request
ip no local secondary nbns address	10	NAS Request
ip no peer ip address	10	NAS Request
ip no peer ip address mask	10	NAS Request
ip no peer primary dns address	10	NAS Request
ip no peer primary nbns address	10	NAS Request
ip no peer secondary dns address	10	NAS Request
ip no peer secondary nbns address	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
ip no service	10	NAS Request
ip peer renegotiate rx conf ack	10	NAS Request
ip peer renegotiate rx conf nak	10	NAS Request
ip peer renegotiate rx conf rej	10	NAS Request
ip peer renegotiate rx conf req	10	NAS Request
ip peer terminate term ack	10	NAS Request
ip peer terminate code rej	10	NAS Request
ip peer terminate term req	10	NAS Request
ip service disable	10	NAS Request
ip stale stacking	10	NAS Request
ipv6 admin disable	10	NAS Request
ipv6 inhibited by authentication	10	NAS Request
ipv6 link down	10	NAS Request
ipv6 local and peer interface ids identical	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
ipv6 max configure exceeded	10	NAS Request
ipv6 no local ipv6 interface id	10	NAS Request
ipv6 no peer ipv6 interface id	10	NAS Request
ipv6 no service	10	NAS Request
ipv6 peer renegotiate rx conf ack	10	NAS Request
ipv6 peer renegotiate rx conf nak	10	NAS Request
ipv6 peer renegotiate rx conf rej	10	NAS Request
ipv6 peer renegotiate rx conf req	10	NAS Request
ipv6 peer terminate code rej	10	NAS Request
ipv6 peer terminate term ack	10	NAS Request
ipv6 peer terminate term req	10	NAS Request
ipv6 service disable	10	NAS Request
ipv6 stale stacking	10	NAS Request
lcp authenticate terminate hold	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
lcp configured mrru too small	10	NAS Request
lcp configured mru invalid	10	NAS Request
lcp configured mru too small	10	NAS Request
lcp dynamic interface hold	10	NAS Request
lcp keepalive failure	10	NAS Request
lcp loopback rx conf req	10	NAS Request
lcp loopback rx echo reply	10	NAS Request
lcp loopback rx echo req	10	NAS Request
lcp max configure exceeded	10	NAS Request
lcp mru changed	10	NAS Request
lcp negotiation timeout	10	NAS Request
lcp no localacm	10	NAS Request
lcp no localacfc	10	NAS Request
lcp no local authentication	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
lcp no local endpoint discriminator	10	NAS Request
lcp no local magic number	10	NAS Request
lcp no local mrru	10	NAS Request
lcp no local mru	10	NAS Request
lcp no local pfc	10	NAS Request
lcp no peer accm	10	NAS Request
lcp no peer authentication	10	NAS Request
lcp no peer endpoint discriminator	10	NAS Request
lcp no peer magicnumber	10	NAS Request
lcp no peer mrru	10	NAS Request
lcp no peer mru	10	NAS Request
lcp no peer pfc	10	NAS Request
lcp peer terminate code rej	1	User Request
lcp peer terminate term ack	1	User Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
lcp peer terminate term req	1	User Request
lcp peer terminate protocol reject	1	User Request
lcp peer renegotiate rx conf ack	1	User Request
lcp peer renegotiate rx conf nak	1	User Request
lcp peer renegotiate rx conf rej	1	User Request
lcp peer renegotiate rx conf req	1	User Request
lcp tunnel disconnected	10	NAS Request
lcp tunnel failed	10	NAS Request
link interface no hardware	8	Port Error
lower interface attach failed	2	Lost Carrier
lower interface teardown	2	Lost Carrier
mpls admin disable	10	NAS Request
mpls link down	10	NAS Request
mpls max configure exceeded	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
mpls no service	10	NAS Request
mpls peer renegotiate rx conf ack	10	NAS Request
mpls peer renegotiate rx conf nak	10	NAS Request
mpls peer renegotiate rx conf rej	10	NAS Request
mpls peer renegotiate rx conf req	10	NAS Request
mpls peer terminate code rej	10	NAS Request
mpls peer terminate term ack	10	NAS Request
mpls peer terminate term req	10	NAS Request
mpls service disable	10	NAS Request
mpls stale stacking	10	NAS Request
network interface admin disable	6	Admin Reset
no bundle	10	NAS Request
no interface	8	Port Error
no link interface	8	Port Error

Table 24: Default Mapping Between PPP Termination Causes and Code Values (Continued)

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
no ncps available	10	NAS Request
no network interface	10	NAS Request
no upper interface	9	NAS Error
osi admin disable	10	NAS Request
osi link down	10	NAS Request
osi max configure exceeded	10	NAS Request
osi no local align npdu	10	NAS Request
osi no peer align npdu	10	NAS Request
osi no service	10	NAS Request
osi peer renegotiate rx conf ack	10	NAS Request
osi peer renegotiate rx conf nak	10	NAS Request
osi peer renegotiate rx conf rej	10	NAS Request
osi peer renegotiate rx conf req	10	NAS Request
osi peer terminate code rej	10	NAS Request

Table 24: Default Mapping Between PPP Termination Causes and Code Values *(Continued)*

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
osi peer terminate term ack	10	NAS Request
osi peer terminate term req	10	NAS Request
osi service disable	10	NAS Request
osi stale stacking	10	NAS Request
recovery active state cleanup	9	NAS Error
recovery configured state cleanup	9	NAS Error
recovery init state cleanup	9	NAS Error
recovery terminated state cleanup	9	NAS Error
recovery terminating state cleanup	9	NAS Error
session init failed	9	NAS Error
subscriber mgr activation failed	9	NAS Error
subscriber mgr get credentials failed	9	NAS Error
subscriber mgr link interface not found	9	NAS Error
subscriber mgr set state active failed	9	NAS Error

RELATED DOCUMENTATION

| [Session Termination Causes and RADIUS Termination Cause Codes](#) | 774

VLAN Termination Causes and Code Values

When a VLAN event terminates a subscriber session, causing a RADIUS Acct-Stop message to be issued, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting*, define the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes that are mapped to the RFC-defined code values. When a subscriber session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot terminated sessions.

[Table 25 on page 822](#) lists the default mapping between the internal identifier for VLAN termination causes and the code values that represent them in the RADIUS Acct-Terminate-Cause attribute (49).



NOTE: You can remap the internal identifiers to a custom code value in the range 1 through 4,294,967,295 by using the `terminate-code` statement at the `[edit access]` hierarchy level. You can view the current mapping by issuing the `show network-access terminate-code vlan detail` command.

Table 25: Default Mapping Between VLAN Termination Causes and Code Values

Internal VLAN Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
vlan admin-logout	6	VLAN session termination initiated by the subscriber being administratively logged out.
vlan admin-reconnect	16	VLAN session termination initiated by the subscriber being administratively reconnected.

Table 25: Default Mapping Between VLAN Termination Causes and Code Values (Continued)

Internal VLAN Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
vlan other	9	VLAN session termination initiated by an otherwise undefined cause.
vlan out-of-band-access-interface-down	2	VLAN out-of-band session termination initiated by the access-facing interface going down.
vlan out-of-band-admin-access-interface-down	6	VLAN out-of-band session termination initiated by the access-facing interface being brought down administratively.
vlan out-of-band-admin-core-interface-down	6	VLAN out-of-band session termination initiated by the core-facing interface being brought down administratively.
vlan out-of-band-ancp-port-down	1	VLAN out-of-band session termination initiated by the receipt of an ANCP Port Down message.
vlan out-of-band-ancp-port-vlan-id-change	1	VLAN out-of-band session termination initiated by a change in the port VLAN ID, which is treated as a logical Port Down message.
vlan out-of-band-core-interface-down	2	VLAN out-of-band session termination initiated by the core-facing interface going down.
vlan out-of-band-l2-wholesale-no-free-vlans	15	VLAN out-of-band session termination initiated by the lack of any remaining available VLAN IDs.
vlan profile-request-error	9	VLAN session termination initiated by an error when requesting the dynamic profile associated with the VLAN range.
vlan sdb-error	9	VLAN session termination initiated by an error in the session database.

Table 25: Default Mapping Between VLAN Termination Causes and Code Values *(Continued)*

Internal VLAN Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
vlan subscriber-activate-error	9	VLAN session termination initiated by an error while attempting to activate the subscriber services for the session.

RELATED DOCUMENTATION

| [Session Termination Causes and RADIUS Termination Cause Codes](#) | 774

CHAPTER 3

Domain Maps for Subscriber Management

IN THIS CHAPTER

- Mapping Subscriber Domains to Access and Session Options | 825
- Verifying Domain Maps | 844

Mapping Subscriber Domains to Access and Session Options

IN THIS SECTION

- Domain Mapping Overview | 826
- Configuring a Domain Map | 830
- Configuring a Wildcard Domain Map | 832
- Specifying an Access Profile in a Domain Map | 833
- Specifying an Address Pool in a Domain Map | 834
- Specifying a Dynamic Profile in a Domain Map | 834
- Specifying an AAA Logical System/Routing Instance in a Domain Map | 835
- Specifying a Target Logical System/Routing Instance in a Domain Map | 836
- Specifying a Tunnel Profile in a Domain Map | 837
- Specifying a Tunnel Switch Profile in a Domain Map | 837
- Configuring Domain and Realm Name Usage for Domain Maps | 838
- Specifying Domain and Realm Name Delimiters | 839
- Specifying the Parsing Order for Domain and Realm Names | 840
- Specifying the Parsing Direction for Domain and Realm Names | 840
- Enabling Domain Name Stripping | 842
- Changing the Username and Password to Simplify Off-Chassis Provisioning | 842

Domain Mapping Overview

IN THIS SECTION

- [Types of Domain Maps and Their Order of Precedence | 828](#)
- [Wildcard Domain Map | 828](#)
- [Default Domain Map | 828](#)
- [Domain Map for Subscriber Usernames With No Domain or Realm Name | 829](#)
- [Understanding Domain Maps and Logical System/Routing Instance Contexts | 829](#)
- [Benefits of Using Domain Maps | 830](#)

Domain mapping enables you to configure a map that specifies access options and session-specific parameters. The map is based on the domain name of subscriber sessions — the router applies the mapped options and parameters to sessions for subscribers that have the specified domains. For example, you might configure a domain map that is based on the domain name `example.com`. The options and parameters in that domain map are then applied when subscribers with the specified domain name (for example, `bob@example.com`, `raj@example.com`, and `juan@example.com`) request a AAA service.



NOTE: A subscriber's username is typically made up of two parts — the user's name followed by the user's domain name, which are separated by a delimiter character. The domain name is always to the right of the domain delimiter. For example, in the username, `juan@example.com`, the user's name, `juan` is followed by the domain name `example.com`, and the two are separated by the `@` delimiter character.

However, some systems use a username format in which the domain name *precedes* the user's name. To avoid confusion with the typical domain name usage, this type of preceding domain name is referred to as a realm name, and the realm name is to the left of the realm delimiter. For example, in the username, `top321-example.com/mary`, the `top321-example.com` part is the realm name, `mary` is the user's name, and the `/` character is the delimiter character.

The domain map provides efficiency, and enables you to make changes for a large number of subscribers in one operation. For example, if an address assignment pool becomes exhausted due to the number of subscribers obtaining addresses from the pool, you can create a domain map that specifies that subscribers in a particular domain obtain addresses from a different pool. In another use of the domain map, you might create a new dynamic profile and then configure the domain map to specify which subscribers (by their domain) use that dynamic profile.

Starting in Junos OS Release 21.3R1, you can configure subdomains under a domain map. In a subdomain, you can configure access profiles per VLAN or for a VLAN range. This enhancement gives you the flexibility to differentiate the users in a domain and to provide different services based on the users' profiles.



NOTE: Subscriber management is supported in the default logical system only. The documentation for the subscriber management domain mapping feature describes using the `aaa-logical-system` and `target-logical-system` statements to configure mapping to a non-default logical system. These statements are for future extensions of subscriber management.

Table 26 on page 827 describes the access options and parameters you can configure in the domain map.

Table 26: Domain Map Options and Parameters

Option	Description
AAA logical system/routing instance	Logical system/routing instance in which AAA sends authentication and accounting requests for the subscriber sessions. Subscriber management is supported in the default logical system only.
Access profile	Access profile applied to subscriber sessions.
Address pool	Address pool used to allocate addresses to subscribers.
Domain and realm name rules	Rules for domain and realm name usage, including domain name stripping, supported delimiters, and parse direction (delimiters and the parse direction are configured globally).
Dynamic profile	Dynamic profile applied to subscriber sessions.
PADN parameters	PPPoE route information for subscriber sessions.
Target logical system/routing instance	Logical system/routing instance to which the subscriber interface is attached. Subscriber management is supported in the default logical system only.

Table 26: Domain Map Options and Parameters *(Continued)*

Option	Description
Tunnel profile	Tunnel profile applied to subscriber sessions.

Types of Domain Maps and Their Order of Precedence

Starting in Junos OS Release 16.1, subscriber management uses a specific order when searching for a domain map that matches the subscriber domain name. The following list shows that order:

- Exact match domain map—The subscriber domain name is an exact match to a configured domain map.
- Wildcard domain map—The subscriber domain name is a partial match to a wildcard domain map.
- default domain map—The subscriber domain name is neither an exact match nor a partial wildcard match to a domain map.



NOTE: If the subscriber username does not have a domain name, then no search is performed and the subscriber is associated with the none domain map, if configured.

Wildcard Domain Map

Starting in Junos OS Release 16.1, the wildcard domain map feature enables you to specify a domain name that is used by subscribers when there is no exact match to the subscriber's domain name. For example, if you create a wildcard domain map with the name `xyz*.example.com`, subscribers with the domain names `xyz.example.com`, `xyz-1234.example.com`, `xyz-eastern.example.com`, and `xyz-northern.example.com` are all mapped to that wildcard domain if there was no exact match for the subscribers' domain names. You can insert the asterisk wildcard character anywhere within the domain map to create the desired matching specification. Wildcard domain mapping is also used in cases where subscriber names are derived from the DHCPv4 Agent Remote ID (option 82 suboption 2) or the DHCPv6 Remote-ID (option 37).

Default Domain Map

You can configure a default domain map that the router uses for subscribers whose domain or realm name does not explicitly match any existing domain map, and also is not a partial match to a wildcard domain map. Specify the name `default` as the domain map *domain-map-name*.

For example, you might configure the default domain map to provide limited feature support for guest subscribers, such as a specific address pool used for guests or the routing instance that provides AAA

services. When the router is unable to provide an exact or wildcard match for the guest subscriber, the router then uses the rules specified in the default domain map configuration to handle the guest subscriber's request.

Domain Map for Subscriber Usernames With No Domain or Realm Name

In some cases a subscriber username might not include a domain name or realm name—you can configure a specific domain map that the router uses for these subscribers. Specify the name *none* as the domain map *domain-map-name*.

Understanding Domain Maps and Logical System/Routing Instance Contexts

You can use a domain map to manage the logical system/routing instance that subscriber management uses for AAA and subscriber contexts. Subscriber management is supported in the default logical system only, so you manage the contexts by configuring the routing instance. The following list describes the two types of contexts:

- **Subscriber context**—The logical system/routing instance in which the subscriber interface is placed. For most dynamic subscriber sessions, the initial subscriber session context is the default logical system and default routing instance. One exception is LNS, in which the initial context for a dynamic LNS session (PPP over L2TP) is the same as the peer interface (the LAC facing interface). Therefore, for LNS sessions, if the peer interface uses a non-default routing instance, then the initial context of the subscriber session also uses that non-default routing instance.
- **AAA context**—The logical system/routing instance that the subscriber session uses for RADIUS interactions, such as authentication and accounting requests. By default, the AAA context is the same as the initial subscriber context. Therefore, for all subscriber sessions other than dynamic LNS sessions, authentication and authorization is performed in the default logical system/routing instance context, unless the default routing instance is explicitly changed.

You can optionally configure a domain map to use a specific subscriber or AAA context. For example, if a dynamic LNS session is initially created in a non-default routing instance (because the initial subscriber context uses the non-default routing instance), you might use the *target-routing-instance* statement to configure the domain map to place the subscriber in the default routing instance. Or, for security reasons, you might want to have all RADIUS interactions in a particular context. In this case, you would use the *aaa-routing-instance* statement to configure the domain map to change the initial AAA context to the new routing instance.

Using domain maps to manage AAA and subscriber contexts is also useful in layer 3 wholesale environments. For example, you might want to place dynamic VLAN interfaces in different non-default routing instances, while maintaining all RADIUS interactions in the default routing-instance. In this example, the initial AAA context is in the default routing instance, but RADIUS authorization places the subscriber VLAN session in a non-default routing instance. You can then include the *aaa-routing-instance* statement in the domain map, to specify that the AAA context uses the default routing instance for the

dynamic VLAN session. The subscriber session is unchanged and remains in the non-default routing instance.

Benefits of Using Domain Maps

- Domain maps simplify managing subscribers at scale by enabling you to make changes for a large number of subscribers in one operation.
- Domain maps provide granularity in applying changes to specific groups of subscribers based on your map definitions.

Configuring a Domain Map

To configure a domain map for subscriber management:

1. Create the domain map. For the map name, specify the domain name that you want the domain map to use. (Use default for the name of the default domain map.)

```
[edit access]
user@host# edit domain map domain-map-name
```

- For example, to create a domain map to be mapped to subscribers with the domain name example.com:

```
[edit access]
user@host# edit domain map example.com
```

- To create a wildcard domain map to be mapped to subscribers whose domain name is not an exact match, but is a partial match:

```
[edit access]
user@host# edit domain map premiumTier*
```

See ["Configuring a Wildcard Domain Map" on page 832](#).

- To create a default domain map to be mapped to subscribers with non-matching domain names:

```
[edit access]
user@host# edit domain map default
```

- To create a domain map to be mapped to subscribers without a domain or realm name:

```
[edit access]
user@host# edit domain map none
```

2. (Optional) Specify the access profile used to apply access rules for the domain map.
See ["Specifying an Access Profile in a Domain Map" on page 833](#).
3. (Optional) For dynamic profiles, clarify the provided dynamic configuration for the subscriber session.
See ["Specifying a Dynamic Profile in a Domain Map" on page 834](#).
4. (Optional) Specify the address pool used to allocate address for the domain map.
See ["Specifying an Address Pool in a Domain Map" on page 834](#).
5. (Optional) Configure the target logical system/routing instance for the subscriber context.
See ["Specifying an AAA Logical System/Routing Instance in a Domain Map" on page 835](#).
6. (Optional) Configure the target logical system/routing instance in which AAA requests are sent for the domain map.
See ["Specifying a Target Logical System/Routing Instance in a Domain Map" on page 836](#).
7. (Optional) Configure rules for domain names; for example; delimiters, parsing direction, and domain stripping. Delimiters and parsing direction are configured globally for all domain maps. Domain stripping is enabled in the domain map.
See ["Configuring Domain and Realm Name Usage for Domain Maps" on page 838](#).
8. (Optional) Configure rules to remove the domain portion from the username for authentication, accounting, and display purposes.
See ["Enabling Domain Name Stripping" on page 842](#).
9. (Optional) Configure parsing the user portion of the username and strip off the user portion for authentication only.
See ["Changing the Username and Password to Simplify Off-Chassis Provisioning" on page 842](#).
10. (Optional) Specify a password to use for all subscriber authentications for a domain map. This option affects only the username/password sent in the access-request to external policy/RADIUS servers.
See ["Changing the Username and Password to Simplify Off-Chassis Provisioning" on page 842](#).
11. (Optional) Assign a tunnel profile that provides tunnel definitions for the domain map.
See ["Specifying a Tunnel Profile in a Domain Map" on page 837](#).
12. (Optional) Assign a tunnel switch profile to be applied by the domain map.
See ["Specifying a Tunnel Switch Profile in a Domain Map" on page 837](#).

Configuring a Wildcard Domain Map

Subscriber management supports a wildcard domain map feature that enables you to configure a domain mapping that is based on a partial wildcard match. When there is no exact match between the subscriber domain name and a configured domain map, subscriber management next looks for a partial match between the subscriber domain name and a wildcard domain map.

To create the wildcard domain map, you include the asterisk wildcard character when you configure the domain map name, such as, domain map `example*`. You can insert the wildcard character anyplace within the domain map, and the wildcard can represent zero or any number of characters. The asterisk is the only supported wildcard character.

For example, the configuration statement `domain map example*northern.com` creates a wildcard domain map that is a partial match for all domain names beginning with `example` and ending with `northern.com`, such as `examplenorthern.com`, `example-northern.com`, and `example1234northern.com`. However if you move the wildcard character in the domain map name to `domain map example-northern*.com`, this creates a more restrictive match that requires the partial matching domain names to start with `example-northern`, such as `example-northern555.com` or `example-northern-alpha.com`.

Wildcard domain mapping is also useful when subscriber management derives subscriber usernames from the DHCPv4 Agent Remote ID (option 82 suboption 2) or the DHCPv6 Remote-ID (option 37). In these cases, the resultant username is in the format `subscriberID|service-plan|accountID|unused`; for example, `EricSmith|premiumTier1|314159265|0000` (where the `|` character is the delimiter). In this example, subscriber management parses the username left-to-right, and identifies the subscriber's domain as `premiumTier1|314159265|0000`. To create a wildcard domain map that is used for this subscriber, you might configure `domain map premiumTier1*`.

The following example describes how four subscribers are mapped to different domains.

For this example, there are three domain maps configured; the default domain map, a domain map named `example3000.com`, and a wildcard domain map named `example*`. The subscribers are mapped as shown in the following list:

- `eric@example3000.com`—There is an exact domain map match, so the subscriber is mapped to domain `example3000.com`.
- `jack@example1001.com`—There is no exact match, but there is a partial match to the wildcard domain, so the subscriber is mapped to the wildcard domain `example*`.
- `ginger@example-western.com`—There is no exact match, but there is a partial match to the wildcard domain, so the subscriber is also mapped to the wildcard domain `example*`.
- `sunshine@test.com`—There is no exact match, nor is there a partial match to the wildcard domain, so the subscriber is mapped to the default domain.

To configure a wildcard domain map:

1. Specify the domain map name, including the wildcard character.

```
[edit access]
user@host# edit domain map premiumTier*
```

2. Specify the optional characteristics for the wildcard domain map.
See ["Configuring a Domain Map" on page 830](#).

Specifying an Access Profile in a Domain Map

You use access profiles to specify the access rules and options (for example, the RADIUS authentication server and attributes) that the router applies to subscriber sessions. The domain map feature enables you to apply a specific access profile for subscribers in a particular domain.

Access profiles can be specified or modified in several different ways. If conflicts occur, the router applies the access profiles based on the precedence rules shown in [Table 27 on page 833](#).

Table 27: Precedence Rules for Applying Access Profiles

Precedence (High to Low)	How the Access Profile Is Applied
1	Specified by the RADIUS Redirect-VRouter-Name attribute (VSA 26-25)
2	Specified in the domain map configuration stanza
3	Indirectly specified in the domain map configuration stanza by the AAA logical system/routing instance mapping
4	Specified in the client configuration stanza
5	Specified in the logical system/routing instance configuration stanza

To include an access profile in a domain map:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```


2. Specify the access profile you want to include in the domain map.

```
[edit access domain map domain-map-name]
user@host# set access-profile profile-name
```

Specifying an Address Pool in a Domain Map

You can use the domain map feature to specify the address pool that the router uses to allocate address for subscriber sessions. The address pool can include both IPv4 and IPv6 address ranges.

Address pools can be specified or modified in several different ways. If conflicts occur, the router applies the address pool based on the precedence rules shown in [Table 28 on page 834](#).

Table 28: Precedence Rules for Determining the Address Pool to Use

Precedence (High to Low)	How the Address Pool Reference Is Provided
1	Specified by the RADIUS Framed-Pool attribute (RADIUS attribute 88)
2	Configured in the domain map configuration stanza
3	Specified in the client configuration stanza (by address match rules)

To specify the address pool used for a domain map:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the address pool you want to use for the domain map.

```
[edit access domain map domain-map-name]
user@host# set address-pool pool-name
```

Specifying a Dynamic Profile in a Domain Map

A dynamic profile defines the set of characteristics that provide dynamic access and services for subscriber sessions (such as class-of-service, protocols, and interface support). The domain map feature enables you to apply a specific dynamic profile based on subscriber domains.

Dynamic profiles are configured at the [edit dynamic-profiles] hierarchy, and can be specified or modified in several different ways. If conflicts occur, the router applies the dynamic profiles based on the precedence rules shown in [Table 29 on page 835](#).

Table 29: Precedence Rules for Applying Dynamic Profiles

Precedence (High to Low)	How the Dynamic Profile Is Applied
1	Specified by the RADIUS Virtual-Router attribute (VSA 26-1) or the Redirect-VRouter-Name attribute (VSA 26-25)
2	Specified in the domain map configuration stanza
3	Specified in the client configuration stanza

To include a dynamic profile in a domain map:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the dynamic profile you want to include in the domain map.

```
[edit access domain map domain-map-name]
user@host# set dynamic-profile profile-name
```

Specifying an AAA Logical System/Routing Instance in a Domain Map

By default, a domain map uses the subscriber logical system/routing instance as the context in which the authd daemon sends AAA authentication and accounting requests. You can optionally configure the domain map to direct AAA requests to a particular context, based on the subscriber domain name. Specifying a non-default AAA context enables you to manage workflow and traffic load, and to efficiently make changes for a large number of subscribers. For example, after upgrading your RADIUS services, you might configure a domain map to specify that all subscribers in the domain `example.com` are now authenticated by a RADIUS server in a particular AAA context.



NOTE: Changing the AAA context does not change the subscriber context. You use the `target-logical-system` statement to explicitly configure the logical system/routing instance for subscribers.

To configure the logical system/routing instance context used for AAA requests:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the routing instance. If a non-default routing instance is currently configured, you can use the `default` option to specify that the domain map use the default routing instance. The AAA logical system is automatically set to the default.

```
[edit access domain map domain-map-name]
user@host# set aaa-routing-instance (routing-instance-name | default)
```



NOTE: Subscriber management is supported in the default logical system only.

Specifying a Target Logical System/Routing Instance in a Domain Map

By default, the router places a subscriber in the logical system/routing instance context of the interface on which the subscriber negotiations start. You can later change the routing instance of the subscriber's context through the use of either a domain map or the RADIUS authentication server.

Subscriber management is supported in the default logical system only, however you can configure the domain map to use a non-default routing instance. Also, if a non-default routing instance is already configured, you can configure the domain map to use the default routing instance.

To configure the logical system/routing instance context used for a subscriber's interface :

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the target routing instance (the default logical system is used by default). If a non-default routing instance is currently configured, you can use the default option to specify that the domain map use the default routing instance.

```
[edit access domain map domain-map-name]
user@host# set target-routing-instance (routing-instance-name | default)
```



NOTE: Subscriber management is supported in the default logical system only.

Specifying a Tunnel Profile in a Domain Map

Tunnel profiles specify tunnel definitions (for example, a set of L2TP tunnels and their attributes) that the router applies to subscriber sessions. The domain map feature enables you to apply a specific tunnel profile to subscribers in a particular domain.



NOTE: A tunnel profile specified by a RADIUS server in the Tunnel-Group attribute (VSA 26-64) takes precedence over the tunnel profile specified in the domain map.

To include a tunnel profile in a domain map:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the tunnel profile you want to include in the domain map.

```
[edit access domain map domain-map-name]
user@host# set tunnel-profile profile-name
```

SEE ALSO

| *Configuring a Tunnel Profile for Subscriber Access*

Specifying a Tunnel Switch Profile in a Domain Map

Tunnel switch profiles determine whether packets in an L2TP subscriber session from a LAC are switched to another session that has a different destination LNS. The tunnel switch profile can also

specify how certain L2TP AVPs are handled when the packets are switched to a second tunnel. The domain map feature enables you to apply a specific tunnel switch profile to subscribers in a particular domain.



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 domain map. 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.

To include a tunnel switch profile in a domain map:

1. Specify the domain map you want to configure.

```
[edit access]
user@host# edit domain map domain-map-name
```

2. Specify the tunnel switch profile you want to include in the domain map.

```
[edit access domain map domain-map-name]
user@host# set tunnel-switch-profile profile-name
```

SEE ALSO

| *Configuring L2TP Tunnel Switching*

Configuring Domain and Realm Name Usage for Domain Maps

You can configure how the router determines the domain names that are used for the domain mapping feature. At the global level, you can specify rules that are used for domain maps. The global rules enable you to specify additional characters that the router can recognize as domain or realm name delimiters and to specify the direction the router uses to parse domain or realm names. The purpose of parsing a domain or realm name is to identify a single, unique name that the router uses as the subscriber's domain name, regardless of whether the source of the name is in the typical domain name format (joseph@example.com) or in the realm name format (example.com\marilyn). The router uses the resulting domain name for operations such as domain map lookup and processing. At the domain map level, you can also enable domain name stripping. Domain name stripping specifies that the router remove the parsed domain or realm name from the subscriber username prior to performing any additional processing for the domain map.

To configure domain name usage rules for domain maps:

1. (Optional) For domain or realm names, configure the parsing order, which specifies whether the router searches for the domain name or the realm name first.
See ["Specifying the Parsing Order for Domain and Realm Names" on page 840.](#)
2. (Optional) For domain or realm names, configure the delimiters you want the router to recognize for domain maps.
See ["Specifying Domain and Realm Name Delimiters" on page 839.](#)
3. (Optional) For domain or realm names, configure the parse direction you want the router to use when determining domain names for domain maps.
See ["Specifying the Parsing Direction for Domain and Realm Names " on page 840.](#)
4. (Optional) For domain names, configure the router to remove the parsed domain or realm name from usernames in the domain map before using AAA services.
See ["Enabling Domain Name Stripping" on page 842.](#)

Specifying Domain and Realm Name Delimiters

A delimiter is the character that separates a subscriber username from the domain or realm name. Delimiters are commonly used for domain or realm name parsing or domain name stripping. You can specify a maximum of eight delimiters that the router uses to recognize domain or realm names for a domain map. If you do not configure any delimiters, the router uses the @ character by default for domain names. There is no default delimiter for realm names.

For example, your network might include the subscribers bob@test.com, pete!example.com, and test.net\maria. In this case, you would configure the router to recognize the characters @ and ! as domain name delimiters, and the \ character as a realm name delimiter.

Keep the following guidelines in mind when specifying delimiters:

- You cannot use the semicolon (;) as a delimiter.
- If you configure optional domain name delimiters, you must also specify the @ character (the default delimiter) if you want to continue to use it as a delimiter.
- If you configure optional domain name delimiters and then unconfigure them, the router sets the domain map delimiter back to the default @ character.

To configure domain and realm name delimiters for domain maps:

1. Specify that you want to configure domain attributes.

```
[edit]
user@host# edit access domain
```

2. Specify the characters you want to use as domain name delimiters. Do not include spaces between the delimiters.

```
[edit access domain]
user@host# set delimiter @!
```

3. Specify the characters you want to use as realm name delimiters. Do not include spaces between the delimiters.

```
[edit access domain]
user@host# set realm-delimiter \
```

Specifying the Parsing Order for Domain and Realm Names

The router parses the username domain or realm name in order to identify a single, unique name that the router uses as the subscriber's domain name, regardless of whether the source of the name is in the typical domain name format (joseph@example.com) or in the realm name format (example.com\marilyn). You can specify whether the router first searches the subscriber username for a domain name or for a realm name. If the router does not find the specified name (for example, you specify `realm-first` and there is no realm name in the username), then the router searches for the second type of name (domain name, in this case). If the router does not find either a realm-name or a domain name, then there is no domain that can be used for domain mapping operations.

To configure the domain name parsing direction for domain maps:

1. Specify that you want to configure domain attributes.

```
[edit]
user@host# edit access domain
```

2. Specify the parsing order you want the router to use, either the domain name first or the realm name first.

```
[edit access domain]
user@host# set parse-order domain-first
```

Specifying the Parsing Direction for Domain and Realm Names

You can specify the direction in which the router performs the parsing operation it uses to identify subscriber domain or realm names for domain maps. During the parsing operation, the router searches the username until it recognizes a delimiter. It then considers anything to the right of the delimiter as the

domain. By default, the router parses from right to left, starting at the right-most character in the username.

The router uses a subscriber's domain name to perform domain map lookup and processing operations. You can configure how the router identifies a unique domain name when the user's name is presented in a traditional domain name format or a realm name. In the traditional domain name format, the user's name is followed by the domain name; for example, joe@example.com. In the realm name format, the user's name is preceded by the domain name, referred to as the realm name; for example, example.com@joe. The purpose of parsing a domain or realm name is to identify a single name that the router uses as the subscriber's domain name, regardless if the source of the name is the user's original domain name or realm name. The router uses the resulting domain name for operations such as domain map lookup and processing. At the domain map level, you can also enable domain name stripping.

The domain parsing direction you use is important when there are nested domain names. For example, for the username user1@test.com@example.com, right-to-left parsing produces a domain name of example.com. For the same username, left-to-right parsing produces a domain name of test.com@example.com.



NOTE: This operation is similar to parsing the user portion of a username, but the default direction and the results are different.

To configure the domain name parsing direction for domain maps:

1. Specify that you want to configure domain attributes.

```
[edit]
user@host# edit access domain
```

2. Specify the parsing direction you want the router to use if the username uses the typical domain name format, in which the domain name follows the user's name.

```
[edit access domain]
user@host# set parse-direction left-to-right
```

3. Specify the parsing direction you want the router to use if the username uses the realm name format, in which the realm name precedes the user's name.

```
[edit access domain]
user@host# set realm-parse-direction right-to-left
```


Enabling Domain Name Stripping

You can configure the router to strip the domain name from usernames before any AAA services are used. Domain name stripping is done for domain maps. The router uses the delimiters and parsing direction you globally configure to determine the domain name that is removed. For example, if the router uses the default delimiter and parsing direction `right-to-left`, the username `user1@example.com` is stripped to be `user1`.

To configure the router to strip the domain name from usernames in a domain map:

1. Specify the domain map for the stripping operation.

```
[edit]
user@host# edit access domain map domain-map-name
```

2. Enable domain name stripping.

```
[edit access domain map domain-map-name]
user@host# set strip-domain
```

Changing the Username and Password to Simplify Off-Chassis Provisioning

For some use cases, you might want to provision L2TP LAC subscriber usernames and authentication passwords off the router chassis. You can strip off the user portion of the username and override the user password.

You can configure how the router identifies the user portion to be stripped when the username is presented in either the traditional domain name format or the realm name format. In the traditional domain name format, the user's name is followed by the domain name; for example, `joe@example.com`. In the realm name format, the user's name is preceded by the domain name, referred to as the realm name; for example, `example.com@joe`.

You can specify the direction in which the router performs the parsing operation that it uses to identify the user portion of the username. During the parsing operation, the router searches the username until it recognizes a delimiter. By default, the router parses from left to right, starting at the left-most character in the username. Everything to the left of the delimiter is the user portion. This direction works for the traditional domain name format. With this configuration, the router identifies and strips `joe` from `joe@example.com`.

For usernames in the realm name format, you need to change the parsing direction to `right-to-left`. The router parses from right to left, starting at the right-most character in the username. When the router recognizes the delimiter, it considers anything to the right of the delimiter as the user portion. With this configuration, the router identifies and strips `joe` from `example.com@joe`.



NOTE: This operation is similar to the domain name/realm name parsing operation, but the default direction and the results are different than domain name/realm name parsing.



NOTE: The user portion is stripped only for the username sent to an external server for authentication. The unstripped username is used for accounting operations.

To configure the user portion to be stripped from the username for all usernames associated with a domain map:

Specify the parsing direction you want the router to use:

- Use left-to-right when the username is in the typical domain name format, in which the domain name follows the user's name.

```
[edit access domain map domain-map-name]
user@host# set strip-username left-to-right
```

- Use right-to-left when the username is in the realm name format, in which the realm name precedes the user's name.

```
[edit access domain map domain-map-name]
user@host# set strip-username right-to-left
```

You can specify a new password to override the existing password for authenticating any subscriber associated with the domain map. To override the password:

- Specify the override password for PAP authentication.

```
[edit access domain map map-name]
user@host# set override-password password
```

- Specify the override password for CHAP authentication.

```
[edit access domain map map-name]
user@host# set override-chap-password password
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
16.1	Starting in Junos OS Release 16.1, subscriber management uses a specific order when searching for a domain map that matches the subscriber domain name.
16.1	Starting in Junos OS Release 16.1, the wildcard domain map feature enables you to specify a domain name that is used by subscribers when there is no exact match to the subscriber's domain name.

RELATED DOCUMENTATION

| [Verifying Domain Maps](#) | 844

Verifying Domain Maps

IN THIS SECTION

- [Purpose](#) | 844
- [Action](#) | 844

Purpose

Display information related to a domain map.

Action

- To display statistics for the domain map:

```
user@host> show network-access domain-map
```

- To display domain map information for a specific subscriber session:

```
user@host> show network-access aaa subscribers session-id
```

RELATED DOCUMENTATION

[Domain Mapping Overview](#) | 826

[Configuring a Domain Map](#) | 830

Testing and Troubleshooting AAA

IN THIS CHAPTER

- [AAA Testing and Troubleshooting | 846](#)
- [Tracing General Authentication Service \(authd\) Events for Troubleshooting | 854](#)

AAA Testing and Troubleshooting

IN THIS SECTION

- [AAA Configuration Testing and Troubleshooting | 846](#)
- [Testing a Subscriber AAA Configuration | 847](#)

AAA Configuration Testing and Troubleshooting

Subscriber management supports a test feature that enables you to check the AAA configuration of a subscriber. You might use the test feature to verify the subscriber's AAA settings and to help troubleshoot or isolate subscriber login problems. The AAA test process creates a pseudo session that authenticates the subscriber, allocates an address for the subscriber, and issues an accounting start packet. The process then issues an accounting stop request, releases the address, and terminates the pseudo session.

The AAA test results provide details about the attributes that subscriber management assigns to the subscriber during login. The attributes might be assigned by RADIUS, a dynamic profile, static interface configuration, or might be statically assigned. You can test the AAA configuration for DHCP, PPP, and authd-lite subscribers. For L2TP clients, the AAA test process displays all tunnel parameters but does not create an actual tunnel session.



NOTE: The `test aaa` commands support all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see ["RADIUS IETF Attributes Supported by the AAA Service Framework" on page 4](#). For information about Juniper Networks VSAs, see ["Juniper Networks VSAs Supported by the AAA Service Framework" on page 19](#).



NOTE: The `test aaa` commands do not support volume-time accounting (Juniper Networks VSA 26-69 with a value of 2). If volume-time accounting is configured for the test subscriber, the `test` command replaces the statistics with time-only accounting statistics.

Testing a Subscriber AAA Configuration

IN THIS SECTION

● Purpose | 847

● Action | 847

Purpose

Display the AAA attributes that subscriber management assigns to the subscriber during login.

The following example tests the AAA configuration for a PPP subscriber. You can use the `test aaa dhcp user` command to perform a similar test for DHCP subscribers and the `test aaa authd-lite user` command to test authd-lite subscribers.

Action

```
user@host>test aaa ppp user user45@test.net password $ABC123
Authentication Grant
*****User Attributes*****
  User Name -                               user45@test.net
  Client IP Address -                       192.168.1.1
  Client IP Netmask -                       255.255.0.0
  Virtual Router Name -                     default
```

Agent Remote Id -	NULL
Reply Message -	NULL
Primary DNS IP Address -	0.0.0.0
Secondary DNS IP Address -	0.0.0.0
Primary WINS IP Address -	0.0.0.0
Secondary WINS IP Address -	0.0.0.0
Primary DNS IPv6 Address -	::
Secondary DNS IPv6 Address -	::
Framed Pool -	not set
Class Attribute -	TEST
Service Type -	0
Client IPv6 Address -	::
Client IPv6 Mask -	null
Framed IPv6 Prefix -	::/0
Framed IPv6 Pool -	not-set
NDRA IPv6 Prefix -	not-set
Login IPv6 Host -	::
Framed Interface Id -	0:0:0:0
Delegated IPv6 Prefix -	::/0
Delegated IPv6 Pool -	not-set
User Password -	\$ABC123
CHAP Password -	NULL
Mac Address -	00:00:5E:00:53:ab
Idle Timeout -	600
Session Timeout -	6000
Service Name (1) -	cos-service(video_sch, nc_sch)
Service Statistics (1) -	1
Service Acct Interim (1) -	600
Service Activation Type (1) -	1
Service Name (2) -	filter-service(in_filter, out_filter)
Service Statistics (2) -	2
Service Acct Interim (2) -	900
Service Activation Type (2) -	1
Cos shaping rate -	100m
Filter Id -	not set
Framed MTU -	(null)
Framed Route -	not set
Ingress Policy Name -	not set
Egress Policy Name -	not set
IGMP -	disabled
Redirect VR Name -	default
Service Bundle -	Null
Framed Ip Route Tag -	not set

```

Ignore DF Bit - disabled
IGMP Access Group Name - not set
IGMP Access Source Group Name - not set
MLD Access Group Name - not set
MLD Access Source Group Name - not set
IGMP Version - not set
MLD Version - not set
IGMP Immediate Leave - disabled
MLD Immediate Leave - disabled
IPv6 Ingress Policy Name - not set
IPv6 Egress Policy Name - not set
Acct Session ID - 1
Acct Interim Interval - 750
Acct Type - 1
Ingress Statistics - disabled
Egress Statistics - disabled
Chargeable user identity - 0
NAS Port Id - -0/0/0.0
NAS Port - 4095
NAS Port Type - 15
Framed Protocol - 1
IPv4 ADF Rule - 010100
IPv4 ADF Rule - 010101
IPv6 ADF Rule - 030100
IPv6 ADF Rule - 030101
****Pausing 10 seconds before disconnecting the test user*****
Logging out subscriber
    Terminate Id - not set
Test complete. Exiting

```

You can use the agent-remote-id *ari* option with the test aaa dhcp user and test aaa ppp user commands to verify DHCP and PPP subscriber authentication in those networks that support the DSL Forum Agent-Remote-Id (VSA 26-2).

If you specify the DSL Forum Agent-Remote-Id, the output includes the specified value. If you do not specify the VSA, then the Agent-Remote-Id value is shown as NULL.

```

user@host>test aaa ppp user thomastank agent-remote-id "(202)555-1212"

```

```

Authentication Grant

```

```

*****User Attributes*****

```

```

    User Name - thomastank

```



```

Client IP Address - 192.168.1.1
Client IP Netmask - 255.255.0.0
...
NAS Ip Address - 0.0.0.0
Agent Remote Id - (202)555-1212
...

```

The following example shows output when the authentication grant fails due to an invalid password:

```

user@host>test aaa ppp user user45@test.net password 55N33%%56
Authentication Deny
Reason : Access Denied
Received Attributes :
    User Name - user45@test.net
    Client IP Address - 0.0.0.0
    Client IP Netmask - 0.0.0.0
    Virtual Router Name - default
    Agent Remote Id - NULL
    Reply Message - NULL
    Primary DNS IP Address - 0.0.0.0
    Secondary DNS IP Address - 0.0.0.0
    Primary WINS IP Address - 0.0.0.0
    Secondary WINS IP Address - 0.0.0.0
    Primary DNS IPv6 Address - ::
    Secondary DNS IPv6 Address - ::
    Framed Pool - not set
    Class Attribute - not set
    Service Type - 0
    Client IPv6 Address - ::
    Client IPv6 Mask - null
    Framed IPv6 Prefix - ::/0
    Framed IPv6 Pool - not-set
    NDRA IPv6 Prefix - not-set
    Login IPv6 Host - ::
    Framed Interface Id - 0:0:0:0
    Delegated IPv6 Prefix - ::/0
    Delegated IPv6 Pool - not-set
    User Password - 55N33%%56
    CHAP Password - NULL
    Mac Address - 00:00:5E:00:53:ab
    Filter Id - not set
    Framed MTU - (null)

```

```

Framed Route - not set
Ingress Policy Name - not set
Egress Policy Name - not set
IGMP - disabled
Redirect VR Name - default
Service Bundle - Null
Framed Ip Route Tag - not set
Ignore DF Bit - disabled
IGMP Access Group Name - not set
IGMP Access Source Group Name - not set
MLD Access Group Name - not set
MLD Access Source Group Name - not set
IGMP Version - not set
MLD Version - not set
IGMP Immediate Leave - disabled
MLD Immediate Leave - disabled
IPv6 Ingress Policy Name - not set
IPv6 Egress Policy Name - not set
Acct Session ID - 12
Acct Interim Interval - 0
Acct Type - 0
Ingress Statistics - disabled
Egress Statistics - disabled
Chargeable user identity - 0
NAS Port Id - -0/0/0.0
NAS Port - 4095
NAS Port Type - 15
Framed Protocol - 0
Test complete. Exiting

```

For some networks, such as a Layer 2 network with VLAN-OOB subscribers, RADIUS is configured to provide the subscriber address in a client profile with the Client-Profile-Name VSA (26-174). In the default configuration, the test fails when it does not receive a subscriber address directly from RADIUS. To successfully test these subscribers, you must include the `no-address-request` option. The command output displays the client profile name in the Dynamic Profile field and the name of the routing instance conveyed by the Virtual-Router VSA (26-1) in the Routing Instance field.

```
user@host>test aaa ppp user thomastank no-address-request
```

```
Authentication Grant
```

```
*****User Attributes*****
```

```
User Name - thomastank
```

```

Client IP Address -          0.0.0.0
Client IP Netmask -         0.0.0.0
...
IPv6 Egress Policy Name -    not set
Dynamic Profile-            filter-service
Routing Instance -          VR27fin
...

```

Starting in Junos OS Release 19.3R1, the XML output format has changed. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag. The new tag makes it easier to recognize the name/value pairs, both for operators and API clients.



NOTE: You may have to change any scripts that use the XML output to work properly with the new format.

The following example shows an excerpt of sample XML output in the old format:

```

user@host>test aaa ppp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
  <aaa-test-result>
    <aaa-test-status>Authentication Grant</aaa-test-status>
    <aaa-test-status>*****User Attributes*****</aaa-test-status>
    <radius-server-attribute-name>User Name -</radius-server-attribute-name>
    <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
    <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
    <radius-server-attribute-value>default:default</radius-server-attribute-value>
    <radius-server-attribute-name>Service Type -</radius-server-attribute-name>
    <radius-server-attribute-value>Framed</radius-server-attribute-value>
    <radius-server-attribute-name>Agent Remote Id -</radius-server-attribute-name>
    <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
    ...
    <aaa-test-status>Test complete. Exiting</aaa-test-status>
  </aaa-test-result>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>

```

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa ppp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
  <aaa-test-result>
    <aaa-test-status>Authentication Grant</aaa-test-status>
    <aaa-test-status>*****User Attributes*****</aaa-test-status>
    <radius-server-data>
      <radius-server-attribute-name>User Name -</radius-server-attribute-name>
      <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
    </radius-server-data>
    <radius-server-data>
      <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
      <radius-server-attribute-value>default:default</radius-server-attribute-value>
    </radius-server-data>
    <radius-server-data>
      <radius-server-attribute-name>Service Type -</radius-server-attribute-name>
      <radius-server-attribute-value>Framed</radius-server-attribute-value>
    </radius-server-data>
    <radius-server-data>
      <radius-server-attribute-name>Agent Remote Id -</radius-server-attribute-name>
      <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
    </radius-server-data>
    ...
    <aaa-test-status>Test complete. Exiting</aaa-test-status>
  </aaa-test-result>
</cli>
  <banner></banner>
</cli>
</rpc-reply>
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, the XML output format has changed.

RELATED DOCUMENTATION

[AAA Service Framework Overview](#) | 2

Tracing General Authentication Service (authd) Events for Troubleshooting

IN THIS SECTION

- [Configuring the General Authentication Service Trace Log Filename](#) | 855
- [Configuring the Number and Size of General Authentication Service Log Files](#) | 855
- [Configuring Access to the General Authentication Service Log File](#) | 856
- [Configuring a Regular Expression for General Authentication Service Messages to Be Logged](#) | 856
- [Configuring Subscriber Filtering for General Authentication Service Tracing](#) | 857
- [Configuring the General Authentication Service Tracing Flags](#) | 858

The Junos OS trace operations feature tracks general authentication 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. The operations and events are those associated with the authd process, which manages the subscriber AAA infrastructure.

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, `authd`. You can specify a different filename, but you cannot change the directory (`/var/log`) 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](#).

3. 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 general authentication service operations:

Configuring the General Authentication Service Trace Log Filename

By default, the name of the file that records trace output for general authentication service is `authd`. You can specify a different name by including the `file` statement at the `[edit system processes general-authentication-service]` hierarchy level:

To configure the filename for general authentication service tracing operations:

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

```
[edit system processes general-authentication-service traceoptions]
user@host# set file aap_logfile_1
```

Configuring the Number and Size of General Authentication 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, by including the files and size options with the `traceoptions` statement.

```
[edit system processes general-authentication-service traceoptions]
user@host# set file aap_logfile_1 files 20 size 2097152
```

Configuring Access to the General Authentication Service Log File

By default, log files can be accessed only by the user who configures the tracing operation. You can allow 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 general-authentication-service traceoptions]
user@host# set file aap_logfile_1 world-readable
```

To explicitly set the default behavior, in which the log file can only be read by the user who configured tracing:

- Configure the log file to be no-world-readable.

```
[edit system processes general-authentication-service traceoptions]
user@host# set file aap_logfile_1 no-world-readable
```

Configuring a Regular Expression for General Authentication 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 (regex) that will be matched.

To configure regular expressions to match:

- Configure the regular expression.

```
[edit system processes general-authentication-service traceoptions]
user@host# set file aap_logfile_1 match regular-expression
```

Configuring Subscriber Filtering for General Authentication Service Tracing

Starting in Junos OS Release 14.1, you can apply filters to the general authentication 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 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 system processes general-authentication-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 system processes general-authentication-service traceoptions]
user@host# set filter user tom@example.com
```

- Filter results for all subscribers whose username begins with tom.

```
[edit system processes general-authentication-service traceoptions]
user@host# set filter user tom*
```

- Filter results for all subscribers whose username ends with tom.

```
[edit system processes general-authentication-service traceoptions]
user@host# set filter user *tom
```


- Filter results for subscribers with the username tom at all domains beginning with ex.

```
[edit system processes general-authentication-service traceoptions]
user@host# set filter user tom@ex*
```

- Filter results for all subscribers at all domains that end with ample.com.

```
[edit system processes general-authentication-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 system processes general-authentication-service traceoptions]
user@host# set filter user tom*@example.com
```

Configuring the General Authentication 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 system services subscriber-management traceoptions]
user@host# set flag flag
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.1	Starting in Junos OS Release 14.1, you can apply filters to the general authentication service to limit tracing to particular subscribers or domains.

RELATED DOCUMENTATION

| [AAA Service Framework Overview](#) | 2

CHAPTER 5

RADIUS Dictionary Files

IN THIS CHAPTER

- [Junos OS Release 15.1 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 859
- [Junos OS Release 16.1 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 881
- [Junos OS Release 16.2 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 902
- [Junos OS Release 17.1 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 924
- [Junos OS Release 17.4 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 946
- [Junos OS Release 18.2 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 968
- [Junos OS Release 18.4 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 990

Junos OS Release 15.1 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 15.1 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 859

Junos OS Release 15.1 Subscriber Management RADIUS Dictionary [DCT]

```
#####  
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v15.1
```

```

# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 1/13/14).
#
#
# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]

#JUNOS and JUNOS
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string) rt

#JUNOS only
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string) r

#JUNOS only
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string) r

#JUNOS and JUNOS
ATTRIBUTE Unisphere-Primary-Dns             ERX-VSA(4, ipaddr) r

#JUNOS and JUNOS
ATTRIBUTE Unisphere-Secondary-Dns           ERX-VSA(5, ipaddr) r

```

```

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Primary-Wins                ERX-VSA(6, ipaddr)  r

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Secondary-Wins              ERX-VSA(7, ipaddr)  r

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Tunnel-Virtual-Router        ERX-VSA(8, string)  rt

#JUNOS and JUNOSE
# untagged
ATTRIBUTE Unisphere-Tunnel-Password             ERX-VSA(9, string)  rt

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Ingress-Policy-Name          ERX-VSA(10, string) r

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Egress-Policy-Name           ERX-VSA(11, string) r

#JUNOSE only
ATTRIBUTE Unisphere-Ingress-Statistics           ERX-VSA(12, integer) r
VALUE      Unisphere-Ingress-Statistics         disable 0
VALUE      Unisphere-Ingress-Statistics         enable 1

#JUNOSE only
ATTRIBUTE Unisphere-Egress-Statistics            ERX-VSA(13, integer) r
VALUE      Unisphere-Egress-Statistics          disable 0
VALUE      Unisphere-Egress-Statistics          enable 1

#JUNOSE only
ATTRIBUTE Unisphere-Service-Category             ERX-VSA(14, integer) r
VALUE      Unisphere-Service-Category           UBR      1
VALUE      Unisphere-Service-Category           UBRPCR   2
VALUE      Unisphere-Service-Category           nrtVBR 3
VALUE      Unisphere-Service-Category           CBR      4
VALUE      Unisphere-Service-Category           rtVBR   5

#JUNOSE only
ATTRIBUTE Unisphere-pcr                          ERX-VSA(15, integer) r

#JUNOSE only
ATTRIBUTE Unisphere-scr                          ERX-VSA(16, integer) r

```

#JUNOSe only		
ATTRIBUTE Unisphere-mbs	ERX-VSA(17, integer)	r
#JUNOSe only		
ATTRIBUTE Unisphere-Init-CLI-Access-Level	ERX-VSA(18, string)	r
#JUNOSe only		
ATTRIBUTE Unisphere-Allow-All-VR-Access	ERX-VSA(19, integer)	r
VALUE Unisphere-Allow-All-VR-Access	disable	0
VALUE Unisphere-Allow-All-VR-Access	enable	1
#JUNOSe only		
ATTRIBUTE Unisphere-Alt-CLI-Access-Level	ERX-VSA(20, string)	R
#JUNOSe only		
ATTRIBUTE Unisphere-Alt-CLI-VRouter-Name	ERX-VSA(21, string)	R
#JUNOSe only		
ATTRIBUTE Unisphere-SA-Validate	ERX-VSA(22, integer)	r
VALUE Unisphere-SA-Validate	disable	0
VALUE Unisphere-SA-Validate	enable	1
#JUNOS and JUNOSe		
ATTRIBUTE Unisphere-Igmp-enable	ERX-VSA(23, integer)	r
VALUE Unisphere-Igmp-enable	disable	0
VALUE Unisphere-Igmp-enable	enable	1
#JUNOS and JUNOSe		
ATTRIBUTE Unisphere-Pppoe-Description	ERX-VSA(24, string)	
#JUNOS and JUNOSe		
ATTRIBUTE Unisphere-Redirect-VRouter-Name	ERX-VSA(25, string)	rt
#JUNOSe only		
ATTRIBUTE Unisphere-Qos-Profile-Name	ERX-VSA(26, string)	r
#JUNOSe only		
# ERX-VSA(27) reserved		
#JUNOSe only		
ATTRIBUTE Unisphere-Pppoe-Ur1	ERX-VSA(28, string)	r
#JUNOSe only		

```
# ERX-VSA(29) reserved
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Nas-Port-Method ERX-VSA(30, integer) r
```

```
VALUE Unisphere-Nas-Port-Method none 1
```

```
VALUE Unisphere-Nas-Port-Method cisco-avp 2
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Service-Bundle ERX-VSA(31, string) r
```

```
#JUNOSe only
```

```
# ERX-VSA(32) reserved
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Tunnel-Max-Sessions ERX-VSA(33, integer) rt
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag ERX-VSA(34, integer) r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Tunnel-Dialout-Number ERX-VSA(35, string) rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ppp-Username ERX-VSA(36, string) rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ppp-Password ERX-VSA(37, string) rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ppp-Protocol ERX-VSA(38, integer) rt
```

```
VALUE Unisphere-Ppp-Protocol none 0
```

```
VALUE Unisphere-Ppp-Protocol pap 1
```

```
VALUE Unisphere-Ppp-Protocol chap 2
```

```
VALUE Unisphere-Ppp-Protocol pap-chap 3
```

```
VALUE Unisphere-Ppp-Protocol chap-pap 4
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Tunnel-Min-Bps ERX-VSA(39, integer) rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Tunnel-Max-Bps ERX-VSA(40, integer) rt
```

```
#JUNOSe only
```

ATTRIBUTE Unisphere-Tunnel-Bearer-Type	ERX-VSA(41, integer) rt
VALUE Unisphere-Tunnel-Bearer-Type	none 0
VALUE Unisphere-Tunnel-Bearer-Type	analog 1
VALUE Unisphere-Tunnel-Bearer-Type	digital 2
#JUNOS and JUNOSE	
ATTRIBUTE Unisphere-Input-Gigapackets	ERX-VSA(42, integer)
#JUNOS and JUNOSE	
ATTRIBUTE Unisphere-Output-Gigapackets	ERX-VSA(43, integer)
#JUNOSE only	
ATTRIBUTE Unisphere-Tunnel-Interface-Id	ERX-VSA(44, string)
#JUNOSE only	
ATTRIBUTE Unisphere-IpV6-Virtual-Router	ERX-VSA(45, string) r
#JUNOSE only	
ATTRIBUTE Unisphere-IpV6-Local-Interface	ERX-VSA(46, string) r
#JUNOS and JUNOSE	
ATTRIBUTE Unisphere-Ipv6-Primary-Dns	ERX-VSA(47, hexadecimal) r
#JUNOS and JUNOSE	
ATTRIBUTE Unisphere-Ipv6-Secondary-Dns	ERX-VSA(48, hexadecimal) r
#JUNOSE only	
# ERX-VSA(49) reserved	
#JUNOSE only	
# ERX-VSA(50) reserved	
#JUNOSE only	
ATTRIBUTE Unisphere-Disconnect-Cause	ERX-VSA(51, hexadecimal)
#JUNOSE only	
ATTRIBUTE Unisphere-Radius-Client-Address	ERX-VSA(52, ipaddr)
#JUNOSE only	
ATTRIBUTE Unisphere-Service-Description	ERX-VSA(53, string) c
#JUNOSE only	
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size	ERX-VSA(54, integer) rt

```

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Dhcp-Options          ERX-VSA(55, hexadecimal) cRO

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Dhcp-Mac-Addr          ERX-VSA(56, string) c

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Dhcp-Gi-Address        ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-LI-Action              ERX-VSA(58, integer) rt salt-encrypt
VALUE    Unisphere-LI-Action              off    0
VALUE    Unisphere-LI-Action              on     1
VALUE    Unisphere-LI-Action              noop   2

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Med-Dev-Handle         ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Med-Ip-Address         ERX-VSA(60, ipaddr) rt salt-encrypt

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Med-Port-Number        ERX-VSA(61, integer) rt salt-encrypt

#JUNOSE only
ATTRIBUTE Unisphere-MLPPP-Bundle-Name      ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc         ERX-VSA(63, string) c

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Tunnel-Group           ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Deactivate-Service     ERX-VSA(66, string) R

#JUNOS and JUNOSE
ATTRIBUTE Unisphere-Activate-Service-tag1  ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSE only

```



```

ATTRIBUTE Unisphere-Service-Volume-tag1          ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1          ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1            ERX-TAGGED-INT-VSA(69, 1) rt
VALUE      Unisphere-Service-Stats-tag1          disabled          0
VALUE      Unisphere-Service-Stats-tag1          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag1          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1          ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2          ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2           ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2           ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2            ERX-TAGGED-INT-VSA(69, 2) rt
VALUE      Unisphere-Service-Stats-tag2          disabled          0
VALUE      Unisphere-Service-Stats-tag2          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag2          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag2          ERX-TAGGED-INT-VSA(140, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag3          ERX-TAGGED-STRING-VSA(65, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag3           ERX-TAGGED-INT-VSA(67, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag3           ERX-TAGGED-INT-VSA(68, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag3            ERX-TAGGED-INT-VSA(69, 3) rt
VALUE      Unisphere-Service-Stats-tag3          disabled          0
VALUE      Unisphere-Service-Stats-tag3          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag3          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3          ERX-TAGGED-INT-VSA(140, 3) rt

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag4          ERX-TAGGED-STRING-VSA(65, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag4            ERX-TAGGED-INT-VSA(67, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag4            ERX-TAGGED-INT-VSA(68, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag4              ERX-TAGGED-INT-VSA(69, 4) rt
VALUE      Unisphere-Service-Stats-tag4            disabled          0
VALUE      Unisphere-Service-Stats-tag4            timestampOnly     1
VALUE      Unisphere-Service-Stats-tag4            timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4            ERX-TAGGED-INT-VSA(140, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag5           ERX-TAGGED-STRING-VSA(65, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag5             ERX-TAGGED-INT-VSA(67, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag5            ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag5              ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5            disabled          0
VALUE      Unisphere-Service-Stats-tag5            timestampOnly     1
VALUE      Unisphere-Service-Stats-tag5            timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5            ERX-TAGGED-INT-VSA(140, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag6           ERX-TAGGED-STRING-VSA(65, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag6             ERX-TAGGED-INT-VSA(67, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag6            ERX-TAGGED-INT-VSA(68, 6) rt

#JUNOS and JUNOSe

```

```

ATTRIBUTE Unisphere-Service-Stats-tag6      ERX-TAGGED-INT-VSA(69, 6) rt
VALUE     Unisphere-Service-Stats-tag6      disabled          0
VALUE     Unisphere-Service-Stats-tag6      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag6      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6     ERX-TAGGED-INT-VSA(140, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag7    ERX-TAGGED-STRING-VSA(65, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag7     ERX-TAGGED-INT-VSA(67, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag7    ERX-TAGGED-INT-VSA(68, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag7      ERX-TAGGED-INT-VSA(69, 7) rt
VALUE     Unisphere-Service-Stats-tag7      disabled          0
VALUE     Unisphere-Service-Stats-tag7      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag7      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7     ERX-TAGGED-INT-VSA(140, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag8    ERX-TAGGED-STRING-VSA(65, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag8     ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8    ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8      ERX-TAGGED-INT-VSA(69, 8) rt
VALUE     Unisphere-Service-Stats-tag8      disabled          0
VALUE     Unisphere-Service-Stats-tag8      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag8      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8     ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                  ERX-VSA(70, integer) r
VALUE     Unisphere-DF-Bit                  dont-ignore-df-bit 0
VALUE     Unisphere-DF-Bit                  ignore-df-bit 1

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name          ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name       ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name          ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name            ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name        ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name           ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version                ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version              v1      1
VALUE      Unisphere-MLD-Version              v2      2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version               ERX-VSA(78, integer) r
VALUE      Unisphere-IGMP-Version             v1      1
VALUE      Unisphere-IGMP-Version             v2      2
VALUE      Unisphere-IGMP-Version             v3      3

#JUNOSe only
ATTRIBUTE Unisphere-IP-Mcast-Adm-Bw-Limit      ERX-VSA(79, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IPv6-Mcast-Adm-Bw-Limit    ERX-VSA(80, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Access-Loop-Parameters ERX-VSA(81, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Parameters             ERX-VSA(82, string) R0

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Session            ERX-VSA(83, string) r

```

Mobile-IP attributes

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-Algorithm	ERX-VSA(84, integer) r
---	------------------------

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-SPI	ERX-VSA(85, integer) r
-----------------------------------	------------------------

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-Key	ERX-VSA(86, string) r
-----------------------------------	-----------------------

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-Replay	ERX-VSA(87, integer) r
--------------------------------------	------------------------

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-Access-Control	ERX-VSA(88, string) r
--	-----------------------

#JUNOSe only

ATTRIBUTE Unisphere-Mobile-IP-Lifetime	ERX-VSA(89, integer) r
--	------------------------

#JUNOSe only

ATTRIBUTE Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE Unisphere-L2TP-Resynch-Method	disable 0
VALUE Unisphere-L2TP-Resynch-Method	failover 1
VALUE Unisphere-L2TP-Resynch-Method	silent-failover 2
VALUE Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3

#JUNOSe only

ATTRIBUTE Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
---	-----------------------

#JUNOSe only

ATTRIBUTE Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
--	-----------------------

#JUNOSe only

ATTRIBUTE Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
--	-----------------------

#JUNOSe only

ATTRIBUTE Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE Unisphere-Tunnel-Tx-Speed-Method	actual 4

```

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-Query-Interval          ERX-VSA(95, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-Max-Resp-Time           ERX-VSA(96, integer) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Immediate-Leave          ERX-VSA(97, integer) r
VALUE      Unisphere-IGMP-Immediate-Leave         disabled    0
VALUE      Unisphere-IGMP-Immediate-Leave         enabled     1

#JUNOSe only
ATTRIBUTE Unisphere-MLD-Query-Interval           ERX-VSA(98, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-Max-Resp-Time            ERX-VSA(99, integer) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Immediate-Leave           ERX-VSA(100, integer) r
VALUE      Unisphere-MLD-Immediate-Leave          disabled    0
VALUE      Unisphere-MLD-Immediate-Leave          enabled     1

#JUNOSe only
ATTRIBUTE Unisphere-IP-Block-Multicast           ERX-VSA(101, integer) r
VALUE      Unisphere-IP-Block-Multicast          disabled    0
VALUE      Unisphere-IP-Block-Multicast          enabled     1

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-Explicit-Tracking        ERX-VSA(102, integer) r
VALUE      Unisphere-IGMP-Explicit-Tracking      disabled    0
VALUE      Unisphere-IGMP-Explicit-Tracking      enabled     1

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-No-Tracking-V2-Grps     ERX-VSA(103, integer) r
VALUE      Unisphere-IGMP-No-Tracking-V2-Grps    disabled    0
VALUE      Unisphere-IGMP-No-Tracking-V2-Grps    enabled     1

#JUNOSe only
ATTRIBUTE Unisphere-MLD-Explicit-Tracking         ERX-VSA(104, integer) r
VALUE      Unisphere-MLD-Explicit-Tracking       disabled    0
VALUE      Unisphere-MLD-Explicit-Tracking       enabled     1

```

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer) r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled 0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled 1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string) r
-----------	-------------------------------	------------------------

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string) R
-----------	-------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr) R
-----------	-------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
-----------	---------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
-----------	-------------------------------	-----------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
-----------	-------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
-----------	----------------------------	-------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
-----------	----------------------------	-------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c
-----------	----------------------------	-------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer) c
-----------	----------------------------	-------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer) c
-----------	----------------------------	-------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer) c
-----------	----------------------------	-------------------------

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Up          ERX-VSA(119, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Dn          ERX-VSA(120, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up        ERX-VSA(121, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn        ERX-VSA(122, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up       ERX-VSA(123, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up       ERX-VSA(124, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn       ERX-VSA(125, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn       ERX-VSA(126, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Line-State             ERX-VSA(127, integer) c
VALUE      Unisphere-DSL-Line-State           SHOWTIME 1
VALUE      Unisphere-DSL-Line-State           IDLE     2
VALUE      Unisphere-DSL-Line-State           SILENT    3

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Type                   ERX-VSA(128, integer) c
VALUE      Unisphere-DSL-Type                 ADSL1     1
VALUE      Unisphere-DSL-Type                 ADSL2     2
VALUE      Unisphere-DSL-Type                 ADSL2PLUS 3
VALUE      Unisphere-DSL-Type                 VDSL1     4
VALUE      Unisphere-DSL-Type                 VDSL2     5
VALUE      Unisphere-DSL-Type                 SDSL     6
VALUE      Unisphere-DSL-Type                 UNKNOWN  7

#JUNOSe only
ATTRIBUTE Unisphere-Ipv6-NdRa-Prefix           ERX-VSA(129, hexadecimal) r

```



```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Qos-Set-Name          ERX-VSA(130, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DownStream-Calc-Rate   ERX-VSA(141, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-UpStream-Calc-Rate     ERX-VSA(142, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-Max-Clients-Per-Interface   ERX-VSA(143, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-PPP-Ingress-Only       ERX-VSA(144, integer) r
VALUE    Unisphere-PPP-Ingress-Only       disabled  0
VALUE    Unisphere-PPP-Ingress-Only       enabled   1

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-Cos-Scheduler-Pmt-Type      ERX-VSA(146, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Backup-Address-Pool    ERX-VSA(147, string) r

#JUNOSe only
ATTRIBUTE Unisphere-ICR-Partition-Id       ERX-VSA(150, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Octets ERX-VSA(151, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Octets ERX-VSA(152, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Packets ERX-VSA(153, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Packets ERX-VSA(154, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Gigawords ERX-VSA(155, integer)

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Gigawords      ERX-VSA(156, integer)

#JUNOSe only
ATTRIBUTE Jnpr-Ipv6-Ndra-Pool-Name                  ERX-VSA(157, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-PppoE-Padn                            ERX-VSA(158, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Dhcp-Option-82                  ERX-VSA(159, string) c

#JUNOS only
ATTRIBUTE Jnpr-Vlan-Map-Id                          ERX-VSA(160, integer)

#JUNOS only
ATTRIBUTE Jnpr-Ipv6-Delegated-Pool-Name              ERX-VSA(161, string)  r

#JUNOS only
ATTRIBUTE Jnpr-Tx-Connect-Speed                     ERX-VSA(162, integer) r

#JUNOS only
ATTRIBUTE Jnpr-Rx-Connect-Speed                     ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control             ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                           ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port

```

```

# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE      Unisphere-Service-Activate-Type      dynamic-profile 1
VALUE      Unisphere-Service-Activate-Type      op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name          ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate              ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason                  ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords      ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service                 ERX-VSA(180, string) rt

#JUNOS only
# ERX-VSA(181) reserved for 15.1

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id             ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap                 ERX-VSA(183, hexadecimal) c

```

```

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id                ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface            ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map                     ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension               ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload                ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter                ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter               ERX-VSA(192, string) r


#JUNOS and JUNOSe
##### tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1                   ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1             ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1             ERX-TAGGED-INT-VSA(33, 1) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag2                   ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2             ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2             ERX-TAGGED-INT-VSA(33, 2) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag3                   ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag3             ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3             ERX-TAGGED-INT-VSA(33, 3) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag4                   ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag4             ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4             ERX-TAGGED-INT-VSA(33, 4) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag5                   ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5             ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5             ERX-TAGGED-INT-VSA(33, 5) rt

```

ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt

```

ATTRIBUTE Unisphere-Tunnel-Sessions-tag27      ERX-TAGGED-INT-VSA(33, 27) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag28             ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28       ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28       ERX-TAGGED-INT-VSA(33, 28) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag29             ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29       ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29       ERX-TAGGED-INT-VSA(33, 29) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag30             ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30       ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30       ERX-TAGGED-INT-VSA(33, 30) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag31             ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31       ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31       ERX-TAGGED-INT-VSA(33, 31) rt

#JUNOS only
##### tagged service activation type attributes

ATTRIBUTE Unisphere-Service-Act-Type-tag1      ERX-TAGGED-INT-VSA(173, 1) rt
VALUE      Unisphere-Service-Act-Type-tag1     dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag1     op-script 2

ATTRIBUTE Unisphere-Service-Act-Type-tag2      ERX-TAGGED-INT-VSA(173, 2) rt
VALUE      Unisphere-Service-Act-Type-tag2     dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag2     op-script 2

#####
# unisphere.dct - Unisphere ERX dictionary
#####

```

Junos OS Release 16.1 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 16.1 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 881

Junos OS Release 16.1 Subscriber Management RADIUS Dictionary [DCT]

```
#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v16.1
# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 03/18/16).
#
#
# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
```



```
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Virtual-Router ERX-VSA(1, string) rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Address-Pool ERX-VSA(2, string) r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Interface ERX-VSA(3, string) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Dns ERX-VSA(4, ipaddr) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Dns ERX-VSA(5, ipaddr) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Wins ERX-VSA(6, ipaddr) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Wins ERX-VSA(7, ipaddr) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Tunnel-Virtual-Router ERX-VSA(8, string) rt
```

```
#JUNOS and JUNOSe
```

```
# untagged
```

```
ATTRIBUTE Unisphere-Tunnel-Password ERX-VSA(9, string) rt
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Ingress-Policy-Name ERX-VSA(10, string) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Egress-Policy-Name ERX-VSA(11, string) r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ingress-Statistics ERX-VSA(12, integer) r
```

```
VALUE Unisphere-Ingress-Statistics disable 0
```

```
VALUE Unisphere-Ingress-Statistics enable 1
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Egress-Statistics ERX-VSA(13, integer) r
```

VALUE	Unisphere-Egress-Statistics	disable 0
VALUE	Unisphere-Egress-Statistics	enable 1
#JUNOSe only		
ATTRIBUTE	Unisphere-Service-Category	ERX-VSA(14, integer) r
VALUE	Unisphere-Service-Category	UBR 1
VALUE	Unisphere-Service-Category	UBRPCR 2
VALUE	Unisphere-Service-Category	nrtVBR 3
VALUE	Unisphere-Service-Category	CBR 4
VALUE	Unisphere-Service-Category	rtVBR 5
#JUNOSe only		
ATTRIBUTE	Unisphere-pcr	ERX-VSA(15, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-scr	ERX-VSA(16, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-mbs	ERX-VSA(17, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Init-CLI-Access-Level	ERX-VSA(18, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Allow-All-VR-Access	ERX-VSA(19, integer) r
VALUE	Unisphere-Allow-All-VR-Access	disable 0
VALUE	Unisphere-Allow-All-VR-Access	enable 1
#JUNOSe only		
ATTRIBUTE	Unisphere-Alt-CLI-Access-Level	ERX-VSA(20, string) R
#JUNOSe only		
ATTRIBUTE	Unisphere-Alt-CLI-VRouter-Name	ERX-VSA(21, string) R
#JUNOSe only		
ATTRIBUTE	Unisphere-SA-Validate	ERX-VSA(22, integer) r
VALUE	Unisphere-SA-Validate	disable 0
VALUE	Unisphere-SA-Validate	enable 1
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Igmp-enable	ERX-VSA(23, integer) r
VALUE	Unisphere-Igmp-enable	disable 0
VALUE	Unisphere-Igmp-enable	enable 1

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description          ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name      ERX-VSA(25, string)  rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name          ERX-VSA(26, string)  r

#JUNOSe only
# ERX-VSA(27) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Pppoe-UrL                 ERX-VSA(28, string)  r

#JUNOSe only
# ERX-VSA(29) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method           ERX-VSA(30, integer) r
VALUE      Unisphere-Nas-Port-Method         none 1
VALUE      Unisphere-Nas-Port-Method         cisco-avp 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Bundle            ERX-VSA(31, string) r

#JUNOSe only
# ERX-VSA(32) reserved

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Max-Sessions       ERX-VSA(33, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag       ERX-VSA(34, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Dialout-Number     ERX-VSA(35, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Username              ERX-VSA(36, string) rt

#JUNOSe only

```

```

ATTRIBUTE Unisphere-Ppp-Password                ERX-VSA(37, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Protocol                ERX-VSA(38, integer) rt
VALUE      Unisphere-Ppp-Protocol                none      0
VALUE      Unisphere-Ppp-Protocol                pap       1
VALUE      Unisphere-Ppp-Protocol                chap      2
VALUE      Unisphere-Ppp-Protocol                pap-chap  3
VALUE      Unisphere-Ppp-Protocol                chap-pap   4

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Min-Bps              ERX-VSA(39, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Max-Bps              ERX-VSA(40, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Bearer-Type          ERX-VSA(41, integer) rt
VALUE      Unisphere-Tunnel-Bearer-Type          none      0
VALUE      Unisphere-Tunnel-Bearer-Type          analog   1
VALUE      Unisphere-Tunnel-Bearer-Type          digital  2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets            ERX-VSA(42, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets           ERX-VSA(43, integer)

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id          ERX-VSA(44, string)

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router          ERX-VSA(45, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface         ERX-VSA(46, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Primary-Dns             ERX-VSA(47, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Secondary-Dns           ERX-VSA(48, hexadecimal) r

```

```

#JUNOSe only
# ERX-VSA(49) reserved

#JUNOSe only
# ERX-VSA(50) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Disconnect-Cause          ERX-VSA(51, hexadecimal)

#JUNOSe only
ATTRIBUTE Unisphere-Radius-Client-Address      ERX-VSA(52, ipaddr)

#JUNOSe only
ATTRIBUTE Unisphere-Service-Description        ERX-VSA(53, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size      ERX-VSA(54, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Options                ERX-VSA(55, hexadecimal) cRO

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Mac-Addr              ERX-VSA(56, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Gi-Address            ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-LI-Action                  ERX-VSA(58, integer) rt salt-encrypt
VALUE      Unisphere-LI-Action                off    0
VALUE      Unisphere-LI-Action                on     1
VALUE      Unisphere-LI-Action                noop   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Dev-Handle             ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Ip-Address             ERX-VSA(60, ipaddr) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number            ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only

```

```

ATTRIBUTE Unisphere-MLPPP-Bundle-Name          ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc              ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group                ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service          ERX-VSA(66, string) R

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1        ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1         ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1        ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1          ERX-TAGGED-INT-VSA(69, 1) rt
VALUE      Unisphere-Service-Stats-tag1        disabled          0
VALUE      Unisphere-Service-Stats-tag1        timestampOnly      1
VALUE      Unisphere-Service-Stats-tag1        timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1        ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2        ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2         ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2        ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2          ERX-TAGGED-INT-VSA(69, 2) rt
VALUE      Unisphere-Service-Stats-tag2        disabled          0
VALUE      Unisphere-Service-Stats-tag2        timestampOnly      1
VALUE      Unisphere-Service-Stats-tag2        timestampAndVolume 2

```

ATTRIBUTE Unisphere-Service-AcctInt-tag2	ERX-TAGGED-INT-VSA(140, 2) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag3	ERX-TAGGED-STRING-VSA(65, 3) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag3	ERX-TAGGED-INT-VSA(67, 3) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag3	ERX-TAGGED-INT-VSA(68, 3) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Service-Stats-tag3	ERX-TAGGED-INT-VSA(69, 3) rt
VALUE Unisphere-Service-Stats-tag3	disabled 0
VALUE Unisphere-Service-Stats-tag3	timestampOnly 1
VALUE Unisphere-Service-Stats-tag3	timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3	ERX-TAGGED-INT-VSA(140, 3) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag4	ERX-TAGGED-STRING-VSA(65, 4) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag4	ERX-TAGGED-INT-VSA(67, 4) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag4	ERX-TAGGED-INT-VSA(68, 4) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Service-Stats-tag4	ERX-TAGGED-INT-VSA(69, 4) rt
VALUE Unisphere-Service-Stats-tag4	disabled 0
VALUE Unisphere-Service-Stats-tag4	timestampOnly 1
VALUE Unisphere-Service-Stats-tag4	timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4	ERX-TAGGED-INT-VSA(140, 4) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag5	ERX-TAGGED-STRING-VSA(65, 5) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag5	ERX-TAGGED-INT-VSA(67, 5) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag5	ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag5      ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5      disabled          0
VALUE      Unisphere-Service-Stats-tag5      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag5      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5      ERX-TAGGED-INT-VSA(140, 5) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag6      ERX-TAGGED-STRING-VSA(65, 6) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Volume-tag6      ERX-TAGGED-INT-VSA(67, 6) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Timeout-tag6      ERX-TAGGED-INT-VSA(68, 6) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag6      ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6      disabled          0
VALUE      Unisphere-Service-Stats-tag6      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag6      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6      ERX-TAGGED-INT-VSA(140, 6) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag7      ERX-TAGGED-STRING-VSA(65, 7) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Volume-tag7      ERX-TAGGED-INT-VSA(67, 7) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Timeout-tag7      ERX-TAGGED-INT-VSA(68, 7) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag7      ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7      disabled          0
VALUE      Unisphere-Service-Stats-tag7      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag7      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7      ERX-TAGGED-INT-VSA(140, 7) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag8      ERX-TAGGED-STRING-VSA(65, 8) rt
```

#JUNOSe only


```

ATTRIBUTE Unisphere-Service-Volume-tag8          ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8          ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8            ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8          disabled          0
VALUE      Unisphere-Service-Stats-tag8          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag8          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8          ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                        ERX-VSA(70, integer) r
VALUE      Unisphere-DF-Bit                      dont-ignore-df-bit 0
VALUE      Unisphere-DF-Bit                      ignore-df-bit 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name              ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name          ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name            ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name              ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name          ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name            ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version                  ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version                v1 1
VALUE      Unisphere-MLD-Version                v2 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version                 ERX-VSA(78, integer) r
VALUE      Unisphere-IGMP-Version                v1 1

```

VALUE	Unisphere-IGMP-Version	v2 2
VALUE	Unisphere-IGMP-Version	v3 3
#JUNOSe only		
ATTRIBUTE	Unisphere-IP-Mcast-Adm-Bw-Limit	ERX-VSA(79, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-IPv6-Mcast-Adm-Bw-Limit	ERX-VSA(80, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Access-Loop-Parameters	ERX-VSA(81, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Qos-Parameters	ERX-VSA(82, string) R0
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Service-Session	ERX-VSA(83, string) r
# Mobile-IP attributes		
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Algorithm	ERX-VSA(84, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-SPI	ERX-VSA(85, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Key	ERX-VSA(86, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Replay	ERX-VSA(87, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Access-Control	ERX-VSA(88, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Lifetime	ERX-VSA(89, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE	Unisphere-L2TP-Resynch-Method	disable 0
VALUE	Unisphere-L2TP-Resynch-Method	failover 1
VALUE	Unisphere-L2TP-Resynch-Method	silent-failover 2

VALUE	Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE	Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE	Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE	Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE	Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE	Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE	Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE	Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE	Unisphere-Tunnel-Tx-Speed-Method	service-profile 7
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Query-Interval	ERX-VSA(95, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Max-Resp-Time	ERX-VSA(96, integer) r
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-IGMP-Immediate-Leave	ERX-VSA(97, integer) r
VALUE	Unisphere-IGMP-Immediate-Leave	disabled 0
VALUE	Unisphere-IGMP-Immediate-Leave	enabled 1
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Query-Interval	ERX-VSA(98, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Max-Resp-Time	ERX-VSA(99, integer) r
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-MLD-Immediate-Leave	ERX-VSA(100, integer) r
VALUE	Unisphere-MLD-Immediate-Leave	disabled 0
VALUE	Unisphere-MLD-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Block-Multicast	ERX-VSA(101, integer) r
VALUE	Unisphere-IP-Block-Multicast	disabled 0
VALUE	Unisphere-IP-Block-Multicast	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Explicit-Tracking	ERX-VSA(102, integer) r
VALUE	Unisphere-IGMP-Explicit-Tracking	disabled 0
VALUE	Unisphere-IGMP-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-No-Tracking-V2-Grps	ERX-VSA(103, integer) r
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	disabled 0
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Explicit-Tracking	ERX-VSA(104, integer) r
VALUE	Unisphere-MLD-Explicit-Tracking	disabled 0
VALUE	Unisphere-MLD-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer) r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled 0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled 1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string) r
-----------	-------------------------------	------------------------

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string) R
-----------	-------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr) R
-----------	-------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
-----------	---------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
-----------	-------------------------------	-----------------------------

#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Up	ERX-VSA(119, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Dn	ERX-VSA(120, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up	ERX-VSA(121, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn	ERX-VSA(122, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up	ERX-VSA(123, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up	ERX-VSA(124, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn	ERX-VSA(125, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn ERX-VSA(126, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DSL-Line-State ERX-VSA(127, integer) c

VALUE Unisphere-DSL-Line-State SHOWTIME 1

VALUE Unisphere-DSL-Line-State IDLE 2

VALUE Unisphere-DSL-Line-State SILENT 3

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DSL-Type ERX-VSA(128, integer) c

VALUE Unisphere-DSL-Type ADSL1 1

VALUE Unisphere-DSL-Type ADSL2 2

VALUE Unisphere-DSL-Type ADSL2PLUS 3

VALUE Unisphere-DSL-Type VDSL1 4

VALUE Unisphere-DSL-Type VDSL2 5

VALUE Unisphere-DSL-Type SDSL 6

VALUE Unisphere-DSL-Type UNKNOWN 7

#JUNOSe only

ATTRIBUTE Unisphere-Ipv6-NdRa-Prefix ERX-VSA(129, hexadecimal) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-Qos-Set-Name ERX-VSA(130, string) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DownStream-Calc-Rate ERX-VSA(141, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-UpStream-Calc-Rate ERX-VSA(142, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Jnpr-Max-Clients-Per-Interface ERX-VSA(143, integer) r

#JUNOSe only

ATTRIBUTE Unisphere-PPP-Ingress-Only ERX-VSA(144, integer) r

VALUE Unisphere-PPP-Ingress-Only disabled 0

VALUE Unisphere-PPP-Ingress-Only enabled 1

#JUNOS and JUNOSe

ATTRIBUTE Jnpr-Cos-Scheduler-Pmt-Type ERX-VSA(146, string) R


```

ATTRIBUTE Jnpr-Rx-Connect-Speed                ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control        ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                      ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE      Unisphere-Service-Activate-Type    dynamic-profile 1
VALUE      Unisphere-Service-Activate-Type    op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name        ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

```



```

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate          ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason              ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords  ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service            ERX-VSA(180, string) rt

#JUNOS only
# ERX-VSA(181) reserved for 15.1

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id        ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap            ERX-VSA(183, hexadecimal) c

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id         ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface     ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map              ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension        ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload         ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter         ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter        ERX-VSA(192, string) r

#JUNOS only
ATTRIBUTE ERX-Bulk-CoA-Transaction-Id         ERX-VSA(194, integer) r
ATTRIBUTE ERX-Bulk-CoA-Identifier             ERX-VSA(195, integer) r

```

#JUNOS only

ATTRIBUTE Unisphere-IPv4-Input-Service-Set	ERX-VSA(196, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Set	ERX-VSA(197, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Filter	ERX-VSA(198, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Filter	ERX-VSA(199, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Set	ERX-VSA(200, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Set	ERX-VSA(201, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Filter	ERX-VSA(202, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Filter	ERX-VSA(203, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name	ERX-VSA(204, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Rule-Name	ERX-VSA(205, string) R

#JUNOS and JUNOSe

tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1	ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1	ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1	ERX-TAGGED-INT-VSA(33, 1) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag2	ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2	ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2	ERX-TAGGED-INT-VSA(33, 2) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag3	ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag3	ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3	ERX-TAGGED-INT-VSA(33, 3) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag4	ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt

ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt

```

ATTRIBUTE Unisphere-Tunnel-Password-tag29      ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29      ERX-TAGGED-INT-VSA(33, 29) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag30            ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30      ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30      ERX-TAGGED-INT-VSA(33, 30) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag31            ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31      ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31      ERX-TAGGED-INT-VSA(33, 31) rt

```

#JUNOS only

tagged service activation type attributes

```

ATTRIBUTE Unisphere-Service-Act-Type-tag1      ERX-TAGGED-INT-VSA(173, 1) rt
VALUE      Unisphere-Service-Act-Type-tag1      dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag1      op-script 2

ATTRIBUTE Unisphere-Service-Act-Type-tag2      ERX-TAGGED-INT-VSA(173, 2) rt
VALUE      Unisphere-Service-Act-Type-tag2      dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag2      op-script 2

```

```

#####
# unisphere.dct - Unisphere ERX dictionary
#####

```

Junos OS Release 16.2 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 16.2 Subscriber Management RADIUS Dictionary \[DCT\]](#) | 903

Junos OS Release 16.2 Subscriber Management RADIUS Dictionary [DCT]

```
#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v16.2
# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 10/19/15).
#
#
# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]

#JUNOS and JUNOS
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string)  rt

#JUNOS only
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string)  r

#JUNOS only
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string)  r

#JUNOS and JUNOS
```

```

ATTRIBUTE Unisphere-Primary-Dns                ERX-VSA(4, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Secondary-Dns              ERX-VSA(5, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Primary-Wins               ERX-VSA(6, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Secondary-Wins             ERX-VSA(7, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Virtual-Router      ERX-VSA(8, string)  rt

#JUNOS and JUNOSe
# untagged
ATTRIBUTE Unisphere-Tunnel-Password           ERX-VSA(9, string)  rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ingress-Policy-Name        ERX-VSA(10, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Egress-Policy-Name        ERX-VSA(11, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Ingress-Statistics        ERX-VSA(12, integer) r
VALUE    Unisphere-Ingress-Statistics        disable 0
VALUE    Unisphere-Ingress-Statistics        enable 1

#JUNOSe only
ATTRIBUTE Unisphere-Egress-Statistics         ERX-VSA(13, integer) r
VALUE    Unisphere-Egress-Statistics         disable 0
VALUE    Unisphere-Egress-Statistics         enable 1

#JUNOSe only
ATTRIBUTE Unisphere-Service-Category          ERX-VSA(14, integer) r
VALUE    Unisphere-Service-Category          UBR      1
VALUE    Unisphere-Service-Category          UBRPCR   2
VALUE    Unisphere-Service-Category          nrtVBR 3
VALUE    Unisphere-Service-Category          CBR      4
VALUE    Unisphere-Service-Category          rtVBR   5

#JUNOSe only

```

```

ATTRIBUTE Unisphere-pcr                                ERX-VSA(15, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-scr                                ERX-VSA(16, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-mbs                                ERX-VSA(17, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Init-CLI-Access-Level              ERX-VSA(18, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Allow-All-VR-Access                ERX-VSA(19, integer) r
VALUE      Unisphere-Allow-All-VR-Access              disable 0
VALUE      Unisphere-Allow-All-VR-Access              enable 1

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-Access-Level               ERX-VSA(20, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-VRouter-Name              ERX-VSA(21, string) R

#JUNOSe only
ATTRIBUTE Unisphere-SA-Validate                        ERX-VSA(22, integer) r
VALUE      Unisphere-SA-Validate                      disable 0
VALUE      Unisphere-SA-Validate                      enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Igmp-enable                        ERX-VSA(23, integer) r
VALUE      Unisphere-Igmp-enable                      disable 0
VALUE      Unisphere-Igmp-enable                      enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description                  ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name              ERX-VSA(25, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name                  ERX-VSA(26, string) r

#JUNOSe only
# ERX-VSA(27) reserved

```



```
#JUNOSe only
ATTRIBUTE Unisphere-PppE-Ur1          ERX-VSA(28, string)  r
```

```
#JUNOSe only
# ERX-VSA(29) reserved
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method    ERX-VSA(30, integer) r
VALUE    Unisphere-Nas-Port-Method    none 1
VALUE    Unisphere-Nas-Port-Method    cisco-avp 2
```

```
#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Bundle     ERX-VSA(31, string)  r
```

```
#JUNOSe only
# ERX-VSA(32) reserved
```

```
#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Max-Sessions ERX-VSA(33, integer) rt
```

```
#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag ERX-VSA(34, integer) r
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Dialout-Number ERX-VSA(35, string) rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Username        ERX-VSA(36, string)  rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Password        ERX-VSA(37, string)  rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Protocol        ERX-VSA(38, integer) rt
VALUE    Unisphere-Ppp-Protocol        none    0
VALUE    Unisphere-Ppp-Protocol        pap      1
VALUE    Unisphere-Ppp-Protocol        chap     2
VALUE    Unisphere-Ppp-Protocol        pap-chap 3
VALUE    Unisphere-Ppp-Protocol        chap-pap 4
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Min-Bps      ERX-VSA(39, integer) rt
```

```

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Max-Bps                ERX-VSA(40, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Bearer-Type            ERX-VSA(41, integer) rt
VALUE      Unisphere-Tunnel-Bearer-Type          none      0
VALUE      Unisphere-Tunnel-Bearer-Type          analog    1
VALUE      Unisphere-Tunnel-Bearer-Type          digital    2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets              ERX-VSA(42, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets             ERX-VSA(43, integer)

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id            ERX-VSA(44, string)

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router            ERX-VSA(45, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface           ERX-VSA(46, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IpV6-Primary-Dns               ERX-VSA(47, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IpV6-Secondary-Dns             ERX-VSA(48, hexadecimal) r

#JUNOSe only
# ERX-VSA(49) reserved

#JUNOSe only
# ERX-VSA(50) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Disconnect-Cause               ERX-VSA(51, hexadecimal)

#JUNOSe only
ATTRIBUTE Unisphere-Radius-Client-Address          ERX-VSA(52, ipaddr)

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Service-Description          ERX-VSA(53, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size        ERX-VSA(54, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Options                  ERX-VSA(55, hexadecimal) cRO

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Mac-Addr                 ERX-VSA(56, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Gi-Address               ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-LI-Action                     ERX-VSA(58, integer) rt salt-encrypt
VALUE      Unisphere-LI-Action                    off    0
VALUE      Unisphere-LI-Action                    on     1
VALUE      Unisphere-LI-Action                    noop   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Dev-Handle                ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Ip-Address                ERX-VSA(60, ipaddr) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number              ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only
ATTRIBUTE Unisphere-MLPPP-Bundle-Name            ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc               ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group                  ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service           ERX-VSA(66, string) R

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1      ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1        ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1       ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1         ERX-TAGGED-INT-VSA(69, 1) rt
VALUE      Unisphere-Service-Stats-tag1       disabled          0
VALUE      Unisphere-Service-Stats-tag1       timestampOnly    1
VALUE      Unisphere-Service-Stats-tag1       timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1       ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2      ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2        ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2       ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2         ERX-TAGGED-INT-VSA(69, 2) rt
VALUE      Unisphere-Service-Stats-tag2       disabled          0
VALUE      Unisphere-Service-Stats-tag2       timestampOnly    1
VALUE      Unisphere-Service-Stats-tag2       timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag2       ERX-TAGGED-INT-VSA(140, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag3      ERX-TAGGED-STRING-VSA(65, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag3        ERX-TAGGED-INT-VSA(67, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag3       ERX-TAGGED-INT-VSA(68, 3) rt

#JUNOS and JUNOSe

```

```

ATTRIBUTE Unisphere-Service-Stats-tag3      ERX-TAGGED-INT-VSA(69, 3) rt
VALUE     Unisphere-Service-Stats-tag3      disabled          0
VALUE     Unisphere-Service-Stats-tag3      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag3      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3     ERX-TAGGED-INT-VSA(140, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag4    ERX-TAGGED-STRING-VSA(65, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag4      ERX-TAGGED-INT-VSA(67, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag4     ERX-TAGGED-INT-VSA(68, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag4      ERX-TAGGED-INT-VSA(69, 4) rt
VALUE     Unisphere-Service-Stats-tag4      disabled          0
VALUE     Unisphere-Service-Stats-tag4      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag4      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4     ERX-TAGGED-INT-VSA(140, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag5    ERX-TAGGED-STRING-VSA(65, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag5      ERX-TAGGED-INT-VSA(67, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag5     ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag5      ERX-TAGGED-INT-VSA(69, 5) rt
VALUE     Unisphere-Service-Stats-tag5      disabled          0
VALUE     Unisphere-Service-Stats-tag5      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag5      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5     ERX-TAGGED-INT-VSA(140, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag6    ERX-TAGGED-STRING-VSA(65, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag6      ERX-TAGGED-INT-VSA(67, 6) rt

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag6          ERX-TAGGED-INT-VSA(68, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag6            ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6          disabled          0
VALUE      Unisphere-Service-Stats-tag6          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag6          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6          ERX-TAGGED-INT-VSA(140, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag7          ERX-TAGGED-STRING-VSA(65, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag7           ERX-TAGGED-INT-VSA(67, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag7           ERX-TAGGED-INT-VSA(68, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag7            ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7          disabled          0
VALUE      Unisphere-Service-Stats-tag7          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag7          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7          ERX-TAGGED-INT-VSA(140, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag8          ERX-TAGGED-STRING-VSA(65, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag8           ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8           ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8            ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8          disabled          0
VALUE      Unisphere-Service-Stats-tag8          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag8          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8          ERX-TAGGED-INT-VSA(140, 8) rt

```

```

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                ERX-VSA(70, integer) r
VALUE     Unisphere-DF-Bit                dont-ignore-df-bit 0
VALUE     Unisphere-DF-Bit                ignore-df-bit 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name       ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name    ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name       ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name         ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name     ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name        ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version             ERX-VSA(77, integer) r
VALUE     Unisphere-MLD-Version            v1 1
VALUE     Unisphere-MLD-Version            v2 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version            ERX-VSA(78, integer) r
VALUE     Unisphere-IGMP-Version           v1 1
VALUE     Unisphere-IGMP-Version           v2 2
VALUE     Unisphere-IGMP-Version           v3 3

#JUNOSe only
ATTRIBUTE Unisphere-IP-Mcast-Adm-Bw-Limit   ERX-VSA(79, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IPv6-Mcast-Adm-Bw-Limit ERX-VSA(80, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Access-Loop-Parameters ERX-VSA(81, string) r

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Parameters          ERX-VSA(82, string) R0

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Session          ERX-VSA(83, string) r

# Mobile-IP attributes

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Algorithm      ERX-VSA(84, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-SPI            ERX-VSA(85, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Key            ERX-VSA(86, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Replay         ERX-VSA(87, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Access-Control ERX-VSA(88, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Lifetime       ERX-VSA(89, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-L2TP-Resynch-Method      ERX-VSA(90, integer) r
VALUE    Unisphere-L2TP-Resynch-Method      disable 0
VALUE    Unisphere-L2TP-Resynch-Method      failover 1
VALUE    Unisphere-L2TP-Resynch-Method      silent-failover 2
VALUE    Unisphere-L2TP-Resynch-Method      failover-with-silent-backup 3

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Switch-Profile    ERX-VSA(91, string) r

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Up-Stream-Data       ERX-VSA(92, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Down-Stream-Data     ERX-VSA(93, string) c

#JUNOS and JUNOSe

```


ATTRIBUTE	Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE	Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE	Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE	Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE	Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE	Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE	Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE	Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE	Unisphere-Tunnel-Tx-Speed-Method	service-profile 7

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Query-Interval	ERX-VSA(95, integer) r
-----------	-------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Max-Resp-Time	ERX-VSA(96, integer) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-IGMP-Immediate-Leave	ERX-VSA(97, integer) r
VALUE	Unisphere-IGMP-Immediate-Leave	disabled 0
VALUE	Unisphere-IGMP-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Query-Interval	ERX-VSA(98, integer) r
-----------	------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Max-Resp-Time	ERX-VSA(99, integer) r
-----------	-----------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-MLD-Immediate-Leave	ERX-VSA(100, integer) r
VALUE	Unisphere-MLD-Immediate-Leave	disabled 0
VALUE	Unisphere-MLD-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Block-Multicast	ERX-VSA(101, integer) r
VALUE	Unisphere-IP-Block-Multicast	disabled 0
VALUE	Unisphere-IP-Block-Multicast	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Explicit-Tracking	ERX-VSA(102, integer) r
VALUE	Unisphere-IGMP-Explicit-Tracking	disabled 0
VALUE	Unisphere-IGMP-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-No-Tracking-V2-Grps	ERX-VSA(103, integer) r
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	disabled 0
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	enabled 1
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Explicit-Tracking	ERX-VSA(104, integer) r
VALUE	Unisphere-MLD-Explicit-Tracking	disabled 0
VALUE	Unisphere-MLD-Explicit-Tracking	enabled 1
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer) r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled 0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled 1
#JUNOS only		
ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string) r
#JUNOS only		
ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string) r
#JUNOS and JUNOSe		
ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string) R
#JUNOS and JUNOSe		
ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr) R
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-Data-Rate-Dn          ERX-VSA(116, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Att-Data-Rate-Up          ERX-VSA(117, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Att-Data-Rate-Dn          ERX-VSA(118, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Up          ERX-VSA(119, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Dn          ERX-VSA(120, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up        ERX-VSA(121, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn        ERX-VSA(122, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up       ERX-VSA(123, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up       ERX-VSA(124, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn       ERX-VSA(125, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn       ERX-VSA(126, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Line-State             ERX-VSA(127, integer) c
VALUE      Unisphere-DSL-Line-State           SHOWTIME 1
VALUE      Unisphere-DSL-Line-State           IDLE    2
VALUE      Unisphere-DSL-Line-State           SILENT   3

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Type                   ERX-VSA(128, integer) c
VALUE      Unisphere-DSL-Type                 ADSL1    1

```

VALUE	Unisphere-DSL-Type	ADSL2	2
VALUE	Unisphere-DSL-Type	ADSL2PLUS	3
VALUE	Unisphere-DSL-Type	VDSL1	4
VALUE	Unisphere-DSL-Type	VDSL2	5
VALUE	Unisphere-DSL-Type	SDSL	6
VALUE	Unisphere-DSL-Type	UNKNOWN	7
#JUNOSe only			
ATTRIBUTE	Unisphere-Ipv6-NdRa-Prefix	ERX-VSA(129, hexadecimal)	r
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Qos-Set-Name	ERX-VSA(130, string)	r
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-DownStream-Calc-Rate	ERX-VSA(141, integer)	c
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-UpStream-Calc-Rate	ERX-VSA(142, integer)	c
#JUNOS and JUNOSe			
ATTRIBUTE	Jnpr-Max-Clients-Per-Interface	ERX-VSA(143, integer)	r
#JUNOSe only			
ATTRIBUTE	Unisphere-PPP-Ingress-Only	ERX-VSA(144, integer)	r
VALUE	Unisphere-PPP-Ingress-Only	disabled	0
VALUE	Unisphere-PPP-Ingress-Only	enabled	1
#JUNOS and JUNOSe			
ATTRIBUTE	Jnpr-Cos-Scheduler-Pmt-Type	ERX-VSA(146, string)	R
#JUNOSe only			
ATTRIBUTE	Unisphere-Backup-Address-Pool	ERX-VSA(147, string)	r
#JUNOSe only			
ATTRIBUTE	Unisphere-ICR-Partition-Id	ERX-VSA(150, string)	
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Ipv6-Acct-Input-Octets	ERX-VSA(151, integer)	
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Ipv6-Acct-Output-Octets	ERX-VSA(152, integer)	

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Packets      ERX-VSA(153, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Packets     ERX-VSA(154, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Gigawords    ERX-VSA(155, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Gigawords   ERX-VSA(156, integer)

#JUNOSe only
ATTRIBUTE Jnpr-IPv6-Ndra-Pool-Name               ERX-VSA(157, string) r

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-PppoE-Padn                        ERX-VSA(158, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Dhcp-Option-82              ERX-VSA(159, string) c

#JUNOS only
ATTRIBUTE Jnpr-Vlan-Map-Id                      ERX-VSA(160, integer)

#JUNOS only
ATTRIBUTE Jnpr-IPv6-Delegated-Pool-Name         ERX-VSA(161, string) r

#JUNOS only
ATTRIBUTE Jnpr-Tx-Connect-Speed                ERX-VSA(162, integer) r

#JUNOS only
ATTRIBUTE Jnpr-Rx-Connect-Speed                ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control        ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                      ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP

```

```

# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE      Unisphere-Service-Activate-Type      dynamic-profile 1
VALUE      Unisphere-Service-Activate-Type      op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name          ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate              ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason                  ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords       ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service                  ERX-VSA(180, string) rt

```

```

#JUNOS only
# ERX-VSA(181) reserved for 15.1

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id          ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap              ERX-VSA(183, hexadecimal) c

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id           ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface      ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map               ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension         ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload          ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter          ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter         ERX-VSA(192, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Set     ERX-VSA(196, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Set    ERX-VSA(197, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Filter  ERX-VSA(198, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Filter ERX-VSA(199, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Set     ERX-VSA(200, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Set    ERX-VSA(201, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Filter  ERX-VSA(202, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Filter ERX-VSA(203, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name      ERX-VSA(204, string) r

#JUNOS and JUNOSe
##### tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1             ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1       ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1       ERX-TAGGED-INT-VSA(33, 1) rt

```

ATTRIBUTE Unisphere-Tunnel-VR-tag2	ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2	ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2	ERX-TAGGED-INT-VSA(33, 2) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag3	ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag3	ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3	ERX-TAGGED-INT-VSA(33, 3) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag4	ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt

ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt

ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29	ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29	ERX-TAGGED-INT-VSA(33, 29) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag30	ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30	ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30	ERX-TAGGED-INT-VSA(33, 30) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag31	ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31	ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31	ERX-TAGGED-INT-VSA(33, 31) rt
#JUNOS only	
##### tagged service activation type attributes	
ATTRIBUTE Unisphere-Service-Act-Type-tag1	ERX-TAGGED-INT-VSA(173, 1) rt
VALUE Unisphere-Service-Act-Type-tag1	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag1	op-script 2

```

ATTRIBUTE Unisphere-Service-Act-Type-tag2          ERX-TAGGED-INT-VSA(173, 2) rt
VALUE      Unisphere-Service-Act-Type-tag2          dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag2          op-script 2

```

```

#####
# unisphere.dct - Unisphere ERX dictionary
#####

```

Junos OS Release 17.1 Subscriber Management RADIUS Dictionary [DCT]

Junos OS Release 17.1 Subscriber Management RADIUS Dictionary [DCT]

```

#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v17.1
# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 03/18/16).
#
#
# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

```

```
# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]
```

```
# for integers t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string)  rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string)  r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Dns            ERX-VSA(4, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Dns          ERX-VSA(5, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Wins           ERX-VSA(6, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Wins         ERX-VSA(7, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Tunnel-Virtual-Router  ERX-VSA(8, string)  rt
```

```
#JUNOS and JUNOSe
```

```
# untagged
```

```
ATTRIBUTE Unisphere-Tunnel-Password        ERX-VSA(9, string)  rt
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Ingress-Policy-Name    ERX-VSA(10, string) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Egress-Policy-Name     ERX-VSA(11, string) r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ingress-Statistics     ERX-VSA(12, integer) r
```

```
VALUE      Unisphere-Ingress-Statistics   disable 0
```

```

VALUE      Unisphere-Ingress-Statistics      enable  1

#JUNOSe only
ATTRIBUTE  Unisphere-Egress-Statistics      ERX-VSA(13, integer) r
VALUE      Unisphere-Egress-Statistics      disable 0
VALUE      Unisphere-Egress-Statistics      enable  1

#JUNOSe only
ATTRIBUTE  Unisphere-Service-Category      ERX-VSA(14, integer) r
VALUE      Unisphere-Service-Category      UBR      1
VALUE      Unisphere-Service-Category      UBRPCR   2
VALUE      Unisphere-Service-Category      nrtVBR  3
VALUE      Unisphere-Service-Category      CBR      4
VALUE      Unisphere-Service-Category      rtVBR   5

#JUNOSe only
ATTRIBUTE  Unisphere-pcr                    ERX-VSA(15, integer) r

#JUNOSe only
ATTRIBUTE  Unisphere-scr                    ERX-VSA(16, integer) r

#JUNOSe only
ATTRIBUTE  Unisphere-mbs                    ERX-VSA(17, integer) r

#JUNOSe only
ATTRIBUTE  Unisphere-Init-CLI-Access-Level  ERX-VSA(18, string) r

#JUNOSe only
ATTRIBUTE  Unisphere-Allow-All-VR-Access    ERX-VSA(19, integer) r
VALUE      Unisphere-Allow-All-VR-Access    disable 0
VALUE      Unisphere-Allow-All-VR-Access    enable  1

#JUNOSe only
ATTRIBUTE  Unisphere-Alt-CLI-Access-Level    ERX-VSA(20, string) R

#JUNOSe only
ATTRIBUTE  Unisphere-Alt-CLI-VRouter-Name    ERX-VSA(21, string) R

#JUNOSe only
ATTRIBUTE  Unisphere-SA-Validate             ERX-VSA(22, integer) r
VALUE      Unisphere-SA-Validate             disable 0
VALUE      Unisphere-SA-Validate             enable  1

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Igmp-enable          ERX-VSA(23, integer) r
VALUE     Unisphere-Igmp-enable          disable 0
VALUE     Unisphere-Igmp-enable          enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description    ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name ERX-VSA(25, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name     ERX-VSA(26, string) r

#JUNOSe only
# ERX-VSA(27) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Pppoe-Ur1            ERX-VSA(28, string) r

#JUNOSe only
# ERX-VSA(29) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method      ERX-VSA(30, integer) r
VALUE     Unisphere-Nas-Port-Method      none 1
VALUE     Unisphere-Nas-Port-Method      cisco-avp 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Bundle       ERX-VSA(31, string) r

#JUNOSe only
# ERX-VSA(32) reserved

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Max-Sessions   ERX-VSA(33, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag   ERX-VSA(34, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Dialout-Number ERX-VSA(35, string) rt

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Username                ERX-VSA(36, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Password                ERX-VSA(37, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Protocol                ERX-VSA(38, integer) rt
VALUE      Unisphere-Ppp-Protocol              none      0
VALUE      Unisphere-Ppp-Protocol              pap        1
VALUE      Unisphere-Ppp-Protocol              chap        2
VALUE      Unisphere-Ppp-Protocol              pap-chap    3
VALUE      Unisphere-Ppp-Protocol              chap-pap     4

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Min-Bps              ERX-VSA(39, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Max-Bps              ERX-VSA(40, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Bearer-Type          ERX-VSA(41, integer) rt
VALUE      Unisphere-Tunnel-Bearer-Type        none      0
VALUE      Unisphere-Tunnel-Bearer-Type        analog    1
VALUE      Unisphere-Tunnel-Bearer-Type        digital   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets           ERX-VSA(42, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets          ERX-VSA(43, integer)

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id         ERX-VSA(44, string)

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router         ERX-VSA(45, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface        ERX-VSA(46, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Primary-Dns            ERX-VSA(47, hexadecimal) r

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Secondary-Dns          ERX-VSA(48, hexadecimal) r

#JUNOSe only
# ERX-VSA(49) reserved

#JUNOSe only
# ERX-VSA(50) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Disconnect-Cause           ERX-VSA(51, hexadecimal)

#JUNOSe only
ATTRIBUTE Unisphere-Radius-Client-Address       ERX-VSA(52, ipaddr)

#JUNOSe only
ATTRIBUTE Unisphere-Service-Description        ERX-VSA(53, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size      ERX-VSA(54, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Options               ERX-VSA(55, hexadecimal) cRO

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Mac-Addr              ERX-VSA(56, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Gi-Address            ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-LI-Action                  ERX-VSA(58, integer) rt salt-encrypt
VALUE      Unisphere-LI-Action                off    0
VALUE      Unisphere-LI-Action                on     1
VALUE      Unisphere-LI-Action                noop   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Dev-Handle             ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Ip-Address             ERX-VSA(60, ipaddr) rt salt-encrypt

```



```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number          ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only
ATTRIBUTE Unisphere-MLPPP-Bundle-Name        ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc           ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group             ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service       ERX-VSA(66, string) R

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1     ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1      ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1     ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1       ERX-TAGGED-INT-VSA(69, 1) rt
VALUE    Unisphere-Service-Stats-tag1       disabled          0
VALUE    Unisphere-Service-Stats-tag1       timestampOnly      1
VALUE    Unisphere-Service-Stats-tag1       timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1     ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2    ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2      ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2     ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe

```

```

ATTRIBUTE Unisphere-Service-Stats-tag2      ERX-TAGGED-INT-VSA(69, 2) rt
VALUE     Unisphere-Service-Stats-tag2      disabled          0
VALUE     Unisphere-Service-Stats-tag2      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag2      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag2     ERX-TAGGED-INT-VSA(140, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag3    ERX-TAGGED-STRING-VSA(65, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag3     ERX-TAGGED-INT-VSA(67, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag3     ERX-TAGGED-INT-VSA(68, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag3      ERX-TAGGED-INT-VSA(69, 3) rt
VALUE     Unisphere-Service-Stats-tag3      disabled          0
VALUE     Unisphere-Service-Stats-tag3      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag3      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3     ERX-TAGGED-INT-VSA(140, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag4    ERX-TAGGED-STRING-VSA(65, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag4     ERX-TAGGED-INT-VSA(67, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag4     ERX-TAGGED-INT-VSA(68, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag4      ERX-TAGGED-INT-VSA(69, 4) rt
VALUE     Unisphere-Service-Stats-tag4      disabled          0
VALUE     Unisphere-Service-Stats-tag4      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag4      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4     ERX-TAGGED-INT-VSA(140, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag5    ERX-TAGGED-STRING-VSA(65, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag5     ERX-TAGGED-INT-VSA(67, 5) rt

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag5          ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag5            ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5          disabled          0
VALUE      Unisphere-Service-Stats-tag5          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag5          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5          ERX-TAGGED-INT-VSA(140, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag6         ERX-TAGGED-STRING-VSA(65, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag6          ERX-TAGGED-INT-VSA(67, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag6          ERX-TAGGED-INT-VSA(68, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag6            ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6          disabled          0
VALUE      Unisphere-Service-Stats-tag6          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag6          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6          ERX-TAGGED-INT-VSA(140, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag7         ERX-TAGGED-STRING-VSA(65, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag7          ERX-TAGGED-INT-VSA(67, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag7          ERX-TAGGED-INT-VSA(68, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag7            ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7          disabled          0
VALUE      Unisphere-Service-Stats-tag7          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag7          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7          ERX-TAGGED-INT-VSA(140, 7) rt

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag8          ERX-TAGGED-STRING-VSA(65, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag8            ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8            ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8              ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8            disabled          0
VALUE      Unisphere-Service-Stats-tag8            timestampOnly    1
VALUE      Unisphere-Service-Stats-tag8            timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8            ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                          ERX-VSA(70, integer) r
VALUE      Unisphere-DF-Bit                        dont-ignore-df-bit 0
VALUE      Unisphere-DF-Bit                        ignore-df-bit    1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name                ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name            ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name              ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name                ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name            ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name              ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version                    ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version                    v1      1
VALUE      Unisphere-MLD-Version                    v2      2

```

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-IGMP-Version	ERX-VSA(78, integer) r
VALUE	Unisphere-IGMP-Version	v1 1
VALUE	Unisphere-IGMP-Version	v2 2
VALUE	Unisphere-IGMP-Version	v3 3

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Mcast-Adm-Bw-Limit	ERX-VSA(79, integer) r
-----------	---------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-IPv6-Mcast-Adm-Bw-Limit	ERX-VSA(80, integer) r
-----------	-----------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-L2c-Access-Loop-Parameters	ERX-VSA(81, string) r
-----------	--------------------------------------	-----------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Qos-Parameters	ERX-VSA(82, string) R0
-----------	--------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Service-Session	ERX-VSA(83, string) r
-----------	---------------------------	-----------------------

Mobile-IP attributes

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-Algorithm	ERX-VSA(84, integer) r
-----------	-------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-SPI	ERX-VSA(85, integer) r
-----------	-------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-Key	ERX-VSA(86, string) r
-----------	-------------------------	-----------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-Replay	ERX-VSA(87, integer) r
-----------	----------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-Access-Control	ERX-VSA(88, string) r
-----------	------------------------------------	-----------------------

#JUNOSe only

ATTRIBUTE	Unisphere-Mobile-IP-Lifetime	ERX-VSA(89, integer) r
-----------	------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE	Unisphere-L2TP-Resynch-Method	disable 0
VALUE	Unisphere-L2TP-Resynch-Method	failover 1
VALUE	Unisphere-L2TP-Resynch-Method	silent-failover 2
VALUE	Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE	Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE	Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE	Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE	Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE	Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE	Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE	Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE	Unisphere-Tunnel-Tx-Speed-Method	service-profile 7
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Query-Interval	ERX-VSA(95, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Max-Resp-Time	ERX-VSA(96, integer) r
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-IGMP-Immediate-Leave	ERX-VSA(97, integer) r
VALUE	Unisphere-IGMP-Immediate-Leave	disabled 0
VALUE	Unisphere-IGMP-Immediate-Leave	enabled 1
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Query-Interval	ERX-VSA(98, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Max-Resp-Time	ERX-VSA(99, integer) r

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-MLD-Immediate-Leave	ERX-VSA(100, integer) r
VALUE	Unisphere-MLD-Immediate-Leave	disabled 0
VALUE	Unisphere-MLD-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Block-Multicast	ERX-VSA(101, integer) r
VALUE	Unisphere-IP-Block-Multicast	disabled 0
VALUE	Unisphere-IP-Block-Multicast	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Explicit-Tracking	ERX-VSA(102, integer) r
VALUE	Unisphere-IGMP-Explicit-Tracking	disabled 0
VALUE	Unisphere-IGMP-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-No-Tracking-V2-Grps	ERX-VSA(103, integer) r
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	disabled 0
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Explicit-Tracking	ERX-VSA(104, integer) r
VALUE	Unisphere-MLD-Explicit-Tracking	disabled 0
VALUE	Unisphere-MLD-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer) r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled 0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled 1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string) r
-----------	-------------------------------	------------------------

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string) R
-----------	-------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr) R
-----------	-------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Up	ERX-VSA(119, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Dn	ERX-VSA(120, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up	ERX-VSA(121, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn	ERX-VSA(122, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up	ERX-VSA(123, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up	ERX-VSA(124, integer) c


```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn          ERX-VSA(125, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn          ERX-VSA(126, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Line-State                ERX-VSA(127, integer) c
VALUE      Unisphere-DSL-Line-State              SHOWTIME 1
VALUE      Unisphere-DSL-Line-State              IDLE      2
VALUE      Unisphere-DSL-Line-State              SILENT     3

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Type                      ERX-VSA(128, integer) c
VALUE      Unisphere-DSL-Type                    ADSL1      1
VALUE      Unisphere-DSL-Type                    ADSL2      2
VALUE      Unisphere-DSL-Type                    ADSL2PLUS 3
VALUE      Unisphere-DSL-Type                    VDSL1      4
VALUE      Unisphere-DSL-Type                    VDSL2      5
VALUE      Unisphere-DSL-Type                    SDSL       6
VALUE      Unisphere-DSL-Type                    UNKNOWN   7

#JUNOSe only
ATTRIBUTE Unisphere-Ipv6-NdRa-Prefix              ERX-VSA(129, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Qos-Set-Name                  ERX-VSA(130, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DownStream-Calc-Rate          ERX-VSA(141, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-UpStream-Calc-Rate            ERX-VSA(142, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-Max-Clients-Per-Interface          ERX-VSA(143, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-PPP-Ingress-Only              ERX-VSA(144, integer) r
VALUE      Unisphere-PPP-Ingress-Only            disabled  0
VALUE      Unisphere-PPP-Ingress-Only            enabled   1

```

#JUNOS and JUNOSe	
ATTRIBUTE Jnpr-Cos-Scheduler-Pmt-Type	ERX-VSA(146, string) R
#JUNOSe only	
ATTRIBUTE Unisphere-Backup-Address-Pool	ERX-VSA(147, string) r
#JUNOSe only	
ATTRIBUTE Unisphere-ICR-Partition-Id	ERX-VSA(150, string)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Octets	ERX-VSA(151, integer)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Octets	ERX-VSA(152, integer)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Packets	ERX-VSA(153, integer)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Packets	ERX-VSA(154, integer)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Gigawords	ERX-VSA(155, integer)
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Gigawords	ERX-VSA(156, integer)
#JUNOSe only	
ATTRIBUTE Jnpr-IPv6-Ndra-Pool-Name	ERX-VSA(157, string) r
#JUNOS and JUNOSe	
ATTRIBUTE Jnpr-PppoE-Padn	ERX-VSA(158, string) R
#JUNOSe only	
ATTRIBUTE Unisphere-Dhcp-Option-82	ERX-VSA(159, string) c
#JUNOS only	
ATTRIBUTE Jnpr-Vlan-Map-Id	ERX-VSA(160, integer)
#JUNOS only	
ATTRIBUTE Jnpr-IPv6-Delegated-Pool-Name	ERX-VSA(161, string) r

```

#JUNOS only
ATTRIBUTE Jnpr-Tx-Connect-Speed                ERX-VSA(162, integer) r

#JUNOS only
ATTRIBUTE Jnpr-Rx-Connect-Speed                ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control        ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                      ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE      Unisphere-Service-Activate-Type    dynamic-profile 1
VALUE      Unisphere-Service-Activate-Type    op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name        ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

```

```

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate          ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason             ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords  ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service            ERX-VSA(180, string) rt

#JUNOS only
# ERX-VSA(181) reserved for 15.1

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id        ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap            ERX-VSA(183, hexadecimal) c

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id         ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface     ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map              ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension        ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload         ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter         ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter        ERX-VSA(192, string) r

```

#JUNOS only

ATTRIBUTE ERX-Bulk-CoA-Transaction-Id	ERX-VSA(194, integer) r
ATTRIBUTE ERX-Bulk-CoA-Identifier	ERX-VSA(195, integer) r

#JUNOS only

ATTRIBUTE Unisphere-IPv4-Input-Service-Set	ERX-VSA(196, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Set	ERX-VSA(197, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Filter	ERX-VSA(198, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Filter	ERX-VSA(199, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Set	ERX-VSA(200, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Set	ERX-VSA(201, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Filter	ERX-VSA(202, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Filter	ERX-VSA(203, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name	ERX-VSA(204, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Rule-Name	ERX-VSA(205, string) R

#JUNOS and JUNOSe

tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1	ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1	ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1	ERX-TAGGED-INT-VSA(33, 1) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag2	ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2	ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2	ERX-TAGGED-INT-VSA(33, 2) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag3	ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag3	ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3	ERX-TAGGED-INT-VSA(33, 3) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag4	ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt

ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt

ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29	ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29	ERX-TAGGED-INT-VSA(33, 29) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag30	ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30	ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30	ERX-TAGGED-INT-VSA(33, 30) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag31	ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31	ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31	ERX-TAGGED-INT-VSA(33, 31) rt

#JUNOS only

tagged service activation type attributes

ATTRIBUTE Unisphere-Service-Act-Type-tag1	ERX-TAGGED-INT-VSA(173, 1) rt
VALUE Unisphere-Service-Act-Type-tag1	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag1	op-script 2
ATTRIBUTE Unisphere-Service-Act-Type-tag2	ERX-TAGGED-INT-VSA(173, 2) rt
VALUE Unisphere-Service-Act-Type-tag2	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag2	op-script 2

#####

unisphere.dct - Unisphere ERX dictionary

#####

Junos OS Release 17.4 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 17.4 Subscriber Management RADIUS Dictionary \[DCT\] | 946](#)

Junos OS Release 17.4 Subscriber Management RADIUS Dictionary [DCT]

```
#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v17.1
# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 03/18/16).
#
#
# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
```

```
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string)  rt
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string)  r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Dns             ERX-VSA(4, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Dns           ERX-VSA(5, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Primary-Wins            ERX-VSA(6, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Secondary-Wins          ERX-VSA(7, ipaddr)  r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Tunnel-Virtual-Router   ERX-VSA(8, string)  rt
```

```
#JUNOS and JUNOSe
```

```
# untagged
```

```
ATTRIBUTE Unisphere-Tunnel-Password         ERX-VSA(9, string)  rt
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Ingress-Policy-Name     ERX-VSA(10, string) r
```

```
#JUNOS and JUNOSe
```

```
ATTRIBUTE Unisphere-Egress-Policy-Name      ERX-VSA(11, string) r
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Ingress-Statistics      ERX-VSA(12, integer) r
```

```
VALUE    Unisphere-Ingress-Statistics      disable 0
```

```
VALUE    Unisphere-Ingress-Statistics      enable 1
```

```
#JUNOSe only
```

```
ATTRIBUTE Unisphere-Egress-Statistics       ERX-VSA(13, integer) r
```

VALUE	Unisphere-Egress-Statistics	disable 0
VALUE	Unisphere-Egress-Statistics	enable 1
#JUNOSe only		
ATTRIBUTE	Unisphere-Service-Category	ERX-VSA(14, integer) r
VALUE	Unisphere-Service-Category	UBR 1
VALUE	Unisphere-Service-Category	UBRPCR 2
VALUE	Unisphere-Service-Category	nrtVBR 3
VALUE	Unisphere-Service-Category	CBR 4
VALUE	Unisphere-Service-Category	rtVBR 5
#JUNOSe only		
ATTRIBUTE	Unisphere-pcr	ERX-VSA(15, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-scr	ERX-VSA(16, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-mbs	ERX-VSA(17, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Init-CLI-Access-Level	ERX-VSA(18, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Allow-All-VR-Access	ERX-VSA(19, integer) r
VALUE	Unisphere-Allow-All-VR-Access	disable 0
VALUE	Unisphere-Allow-All-VR-Access	enable 1
#JUNOSe only		
ATTRIBUTE	Unisphere-Alt-CLI-Access-Level	ERX-VSA(20, string) R
#JUNOSe only		
ATTRIBUTE	Unisphere-Alt-CLI-VRouter-Name	ERX-VSA(21, string) R
#JUNOSe only		
ATTRIBUTE	Unisphere-SA-Validate	ERX-VSA(22, integer) r
VALUE	Unisphere-SA-Validate	disable 0
VALUE	Unisphere-SA-Validate	enable 1
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Igmp-enable	ERX-VSA(23, integer) r
VALUE	Unisphere-Igmp-enable	disable 0
VALUE	Unisphere-Igmp-enable	enable 1

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description          ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name      ERX-VSA(25, string)  rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name          ERX-VSA(26, string)  r

#JUNOSe only
# ERX-VSA(27) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Pppoe-UrL                 ERX-VSA(28, string)  r

#JUNOSe only
# ERX-VSA(29) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method            ERX-VSA(30, integer) r
VALUE      Unisphere-Nas-Port-Method          none 1
VALUE      Unisphere-Nas-Port-Method          cisco-avp 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Bundle            ERX-VSA(31, string) r

#JUNOSe only
# ERX-VSA(32) reserved

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Max-Sessions        ERX-VSA(33, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag        ERX-VSA(34, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Dialout-Number      ERX-VSA(35, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Username               ERX-VSA(36, string) rt

#JUNOSe only

```

```

ATTRIBUTE Unisphere-Ppp-Password                ERX-VSA(37, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Protocol                ERX-VSA(38, integer) rt
VALUE      Unisphere-Ppp-Protocol                none      0
VALUE      Unisphere-Ppp-Protocol                pap       1
VALUE      Unisphere-Ppp-Protocol                chap      2
VALUE      Unisphere-Ppp-Protocol                pap-chap  3
VALUE      Unisphere-Ppp-Protocol                chap-pap   4

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Min-Bps              ERX-VSA(39, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Max-Bps              ERX-VSA(40, integer) rt

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Bearer-Type          ERX-VSA(41, integer) rt
VALUE      Unisphere-Tunnel-Bearer-Type          none      0
VALUE      Unisphere-Tunnel-Bearer-Type          analog   1
VALUE      Unisphere-Tunnel-Bearer-Type          digital  2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets           ERX-VSA(42, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets          ERX-VSA(43, integer)

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id         ERX-VSA(44, string)

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router         ERX-VSA(45, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface        ERX-VSA(46, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Primary-Dns            ERX-VSA(47, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Secondary-Dns          ERX-VSA(48, hexadecimal) r

```

```

#JUNOSe only
# ERX-VSA(49) reserved

#JUNOSe only
# ERX-VSA(50) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Disconnect-Cause          ERX-VSA(51, hexadecimal)

#JUNOSe only
ATTRIBUTE Unisphere-Radius-Client-Address      ERX-VSA(52, ipaddr)

#JUNOSe only
ATTRIBUTE Unisphere-Service-Description        ERX-VSA(53, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size      ERX-VSA(54, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Options                ERX-VSA(55, hexadecimal) cRO

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Mac-Addr              ERX-VSA(56, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Gi-Address             ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-LI-Action                  ERX-VSA(58, integer) rt salt-encrypt
VALUE      Unisphere-LI-Action                off    0
VALUE      Unisphere-LI-Action                on     1
VALUE      Unisphere-LI-Action                noop   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Dev-Handle              ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Ip-Address              ERX-VSA(60, ipaddr) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number             ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only

```

```

ATTRIBUTE Unisphere-MLPPP-Bundle-Name          ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc              ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group                ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service          ERX-VSA(66, string) R

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1        ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1          ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1         ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1           ERX-TAGGED-INT-VSA(69, 1) rt
VALUE      Unisphere-Service-Stats-tag1         disabled          0
VALUE      Unisphere-Service-Stats-tag1         timestampOnly      1
VALUE      Unisphere-Service-Stats-tag1         timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1         ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2        ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2          ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2         ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2           ERX-TAGGED-INT-VSA(69, 2) rt
VALUE      Unisphere-Service-Stats-tag2         disabled          0
VALUE      Unisphere-Service-Stats-tag2         timestampOnly      1
VALUE      Unisphere-Service-Stats-tag2         timestampAndVolume 2

```

ATTRIBUTE Unisphere-Service-AcctInt-tag2	ERX-TAGGED-INT-VSA(140, 2) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag3	ERX-TAGGED-STRING-VSA(65, 3) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag3	ERX-TAGGED-INT-VSA(67, 3) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag3	ERX-TAGGED-INT-VSA(68, 3) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Service-Stats-tag3	ERX-TAGGED-INT-VSA(69, 3) rt
VALUE Unisphere-Service-Stats-tag3	disabled 0
VALUE Unisphere-Service-Stats-tag3	timestampOnly 1
VALUE Unisphere-Service-Stats-tag3	timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3	ERX-TAGGED-INT-VSA(140, 3) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag4	ERX-TAGGED-STRING-VSA(65, 4) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag4	ERX-TAGGED-INT-VSA(67, 4) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag4	ERX-TAGGED-INT-VSA(68, 4) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Service-Stats-tag4	ERX-TAGGED-INT-VSA(69, 4) rt
VALUE Unisphere-Service-Stats-tag4	disabled 0
VALUE Unisphere-Service-Stats-tag4	timestampOnly 1
VALUE Unisphere-Service-Stats-tag4	timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4	ERX-TAGGED-INT-VSA(140, 4) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Activate-Service-tag5	ERX-TAGGED-STRING-VSA(65, 5) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Volume-tag5	ERX-TAGGED-INT-VSA(67, 5) rt
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Timeout-tag5	ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag5      ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5      disabled          0
VALUE      Unisphere-Service-Stats-tag5      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag5      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5      ERX-TAGGED-INT-VSA(140, 5) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag6    ERX-TAGGED-STRING-VSA(65, 6) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Volume-tag6      ERX-TAGGED-INT-VSA(67, 6) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Timeout-tag6     ERX-TAGGED-INT-VSA(68, 6) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag6      ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6      disabled          0
VALUE      Unisphere-Service-Stats-tag6      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag6      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6      ERX-TAGGED-INT-VSA(140, 6) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag7    ERX-TAGGED-STRING-VSA(65, 7) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Volume-tag7      ERX-TAGGED-INT-VSA(67, 7) rt
```

#JUNOSe only

```
ATTRIBUTE Unisphere-Service-Timeout-tag7     ERX-TAGGED-INT-VSA(68, 7) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Service-Stats-tag7      ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7      disabled          0
VALUE      Unisphere-Service-Stats-tag7      timestampOnly     1
VALUE      Unisphere-Service-Stats-tag7      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7      ERX-TAGGED-INT-VSA(140, 7) rt
```

#JUNOS and JUNOSe

```
ATTRIBUTE Unisphere-Activate-Service-tag8    ERX-TAGGED-STRING-VSA(65, 8) rt
```

#JUNOSe only

```

ATTRIBUTE Unisphere-Service-Volume-tag8          ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8          ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8            ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8          disabled          0
VALUE      Unisphere-Service-Stats-tag8          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag8          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8          ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                        ERX-VSA(70, integer) r
VALUE      Unisphere-DF-Bit                      dont-ignore-df-bit 0
VALUE      Unisphere-DF-Bit                      ignore-df-bit 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name              ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name          ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name            ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name              ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name          ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name            ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version                  ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version                v1 1
VALUE      Unisphere-MLD-Version                v2 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version                  ERX-VSA(78, integer) r
VALUE      Unisphere-IGMP-Version                v1 1

```

VALUE	Unisphere-IGMP-Version	v2 2
VALUE	Unisphere-IGMP-Version	v3 3
#JUNOSe only		
ATTRIBUTE	Unisphere-IP-Mcast-Adm-Bw-Limit	ERX-VSA(79, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-IPv6-Mcast-Adm-Bw-Limit	ERX-VSA(80, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Access-Loop-Parameters	ERX-VSA(81, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Qos-Parameters	ERX-VSA(82, string) R0
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Service-Session	ERX-VSA(83, string) r
# Mobile-IP attributes		
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Algorithm	ERX-VSA(84, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-SPI	ERX-VSA(85, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Key	ERX-VSA(86, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Replay	ERX-VSA(87, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Access-Control	ERX-VSA(88, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Mobile-IP-Lifetime	ERX-VSA(89, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE	Unisphere-L2TP-Resynch-Method	disable 0
VALUE	Unisphere-L2TP-Resynch-Method	failover 1
VALUE	Unisphere-L2TP-Resynch-Method	silent-failover 2

VALUE	Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
#JUNOSe only		
ATTRIBUTE	Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE	Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE	Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE	Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE	Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE	Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE	Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE	Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE	Unisphere-Tunnel-Tx-Speed-Method	service-profile 7
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Query-Interval	ERX-VSA(95, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-IGMP-Max-Resp-Time	ERX-VSA(96, integer) r
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-IGMP-Immediate-Leave	ERX-VSA(97, integer) r
VALUE	Unisphere-IGMP-Immediate-Leave	disabled 0
VALUE	Unisphere-IGMP-Immediate-Leave	enabled 1
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Query-Interval	ERX-VSA(98, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-MLD-Max-Resp-Time	ERX-VSA(99, integer) r
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-MLD-Immediate-Leave	ERX-VSA(100, integer) r
VALUE	Unisphere-MLD-Immediate-Leave	disabled 0
VALUE	Unisphere-MLD-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Block-Multicast	ERX-VSA(101, integer) r
VALUE	Unisphere-IP-Block-Multicast	disabled 0
VALUE	Unisphere-IP-Block-Multicast	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Explicit-Tracking	ERX-VSA(102, integer) r
VALUE	Unisphere-IGMP-Explicit-Tracking	disabled 0
VALUE	Unisphere-IGMP-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-No-Tracking-V2-Grps	ERX-VSA(103, integer) r
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	disabled 0
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Explicit-Tracking	ERX-VSA(104, integer) r
VALUE	Unisphere-MLD-Explicit-Tracking	disabled 0
VALUE	Unisphere-MLD-Explicit-Tracking	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer) r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled 0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled 1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string) r
-----------	-------------------------------	------------------------

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string) R
-----------	-------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr) R
-----------	-------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
-----------	---------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
-----------	-------------------------------	-----------------------------

#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Up	ERX-VSA(119, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Dn	ERX-VSA(120, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up	ERX-VSA(121, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn	ERX-VSA(122, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up	ERX-VSA(123, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up	ERX-VSA(124, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn	ERX-VSA(125, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn ERX-VSA(126, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DSL-Line-State ERX-VSA(127, integer) c

VALUE Unisphere-DSL-Line-State SHOWTIME 1

VALUE Unisphere-DSL-Line-State IDLE 2

VALUE Unisphere-DSL-Line-State SILENT 3

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DSL-Type ERX-VSA(128, integer) c

VALUE Unisphere-DSL-Type ADSL1 1

VALUE Unisphere-DSL-Type ADSL2 2

VALUE Unisphere-DSL-Type ADSL2PLUS 3

VALUE Unisphere-DSL-Type VDSL1 4

VALUE Unisphere-DSL-Type VDSL2 5

VALUE Unisphere-DSL-Type SDSL 6

VALUE Unisphere-DSL-Type UNKNOWN 7

#JUNOSe only

ATTRIBUTE Unisphere-Ipv6-NdRa-Prefix ERX-VSA(129, hexadecimal) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-Qos-Set-Name ERX-VSA(130, string) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-DownStream-Calc-Rate ERX-VSA(141, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-UpStream-Calc-Rate ERX-VSA(142, integer) c

#JUNOS and JUNOSe

ATTRIBUTE Jnpr-Max-Clients-Per-Interface ERX-VSA(143, integer) r

#JUNOSe only

ATTRIBUTE Unisphere-PPP-Ingress-Only ERX-VSA(144, integer) r

VALUE Unisphere-PPP-Ingress-Only disabled 0

VALUE Unisphere-PPP-Ingress-Only enabled 1

#JUNOS and JUNOSe

ATTRIBUTE Jnpr-Cos-Scheduler-Pmt-Type ERX-VSA(146, string) R


```

ATTRIBUTE Jnpr-Rx-Connect-Speed                ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control        ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                      ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE      Unisphere-Service-Activate-Type    dynamic-profile 1
VALUE      Unisphere-Service-Activate-Type    op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name        ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

```

```

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate          ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason             ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords  ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service            ERX-VSA(180, string) rt

#JUNOS only
# ERX-VSA(181) reserved for 15.1

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id        ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap            ERX-VSA(183, hexadecimal) c

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id         ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface     ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map             ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension        ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload         ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter         ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter        ERX-VSA(192, string) r

#JUNOS only
ATTRIBUTE ERX-Bulk-CoA-Transaction-Id         ERX-VSA(194, integer) r
ATTRIBUTE ERX-Bulk-CoA-Identifier             ERX-VSA(195, integer) r

```

#JUNOS only

ATTRIBUTE Unisphere-IPv4-Input-Service-Set	ERX-VSA(196, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Set	ERX-VSA(197, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Filter	ERX-VSA(198, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Filter	ERX-VSA(199, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Set	ERX-VSA(200, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Set	ERX-VSA(201, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Filter	ERX-VSA(202, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Filter	ERX-VSA(203, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name	ERX-VSA(204, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Rule-Name	ERX-VSA(205, string) R

#JUNOS only

ATTRIBUTE ERX-DHCPv6-Options	ERX-VSA(207, hexadecimal) CRO
ATTRIBUTE ERX-DHCP-Header	ERX-VSA(208, hexadecimal) c
ATTRIBUTE ERX-DHCPv6-Header	ERX-VSA(209, hexadecimal) c
ATTRIBUTE ERX-Acct-Request-Reason	ERX-VSA(210, integer) c
VALUE ERX-Acct-Request-Reason	Acct-Start-Ack 1
VALUE ERX-Acct-Request-Reason	Periodic-Interval-Interim 2
VALUE ERX-Acct-Request-Reason	IP-Active 4
VALUE ERX-Acct-Request-Reason	IP-Inactive 8
VALUE ERX-Acct-Request-Reason	IPv6-Active 16
VALUE ERX-Acct-Request-Reason	IPv6-Inactive 32
VALUE ERX-Acct-Request-Reason	Session-Active 64
VALUE ERX-Acct-Request-Reason	Session-Inactive 128
VALUE ERX-Acct-Request-Reason	Line-Speed-Change 256
VALUE ERX-Acct-Request-Reason	Address-Assignment-Change 512
VALUE ERX-Acct-Request-Reason	CoA-Processing-Complete 1024

#JUNOS and JUNOSe

tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1	ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1	ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1	ERX-TAGGED-INT-VSA(33, 1) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag2	ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2	ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2	ERX-TAGGED-INT-VSA(33, 2) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag3	ERX-TAGGED-STRING-VSA(8, 3) rt

ATTRIBUTE Unisphere-Tunnel-Password-tag3	ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3	ERX-TAGGED-INT-VSA(33, 3) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag4	ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29	ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29	ERX-TAGGED-INT-VSA(33, 29) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag30	ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30	ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30	ERX-TAGGED-INT-VSA(33, 30) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag31	ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31	ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31	ERX-TAGGED-INT-VSA(33, 31) rt

#JUNOS only

tagged service activation type attributes

ATTRIBUTE Unisphere-Service-Act-Type-tag1	ERX-TAGGED-INT-VSA(173, 1) rt
VALUE Unisphere-Service-Act-Type-tag1	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag1	op-script 2
ATTRIBUTE Unisphere-Service-Act-Type-tag2	ERX-TAGGED-INT-VSA(173, 2) rt
VALUE Unisphere-Service-Act-Type-tag2	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag2	op-script 2

#####

```
# unisphere.dct - Unisphere ERX dictionary
#####
```

Junos OS Release 18.2 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 18.2 Subscriber Management RADIUS Dictionary \[DCT\] | 968](#)

Junos OS Release 18.2 Subscriber Management RADIUS Dictionary [DCT]

```
#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v17.1
# (Note: JUNOS attribute breakdown by JUNOS releases not reflected in this document.
#       JUNOS information is based on complete JUNOS attribute listing as of 03/18/16).
#
#
# (See README.DCT for more details on the format of this file)
#
# please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]
```

```

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string)  rt

#JUNOSe only
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string)  r

#JUNOSe only
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Primary-Dns            ERX-VSA(4, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Secondary-Dns          ERX-VSA(5, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Primary-Wins           ERX-VSA(6, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Secondary-Wins         ERX-VSA(7, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Virtual-Router  ERX-VSA(8, string)  rt

#JUNOS and JUNOSe
# untagged
ATTRIBUTE Unisphere-Tunnel-Password        ERX-VSA(9, string)  rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ingress-Policy-Name    ERX-VSA(10, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Egress-Policy-Name     ERX-VSA(11, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Ingress-Statistics     ERX-VSA(12, integer) r

```



```

VALUE    Unisphere-Ingress-Statistics    disable 0
VALUE    Unisphere-Ingress-Statistics    enable  1

#JUNOSe only
ATTRIBUTE Unisphere-Egress-Statistics    ERX-VSA(13, integer) r
VALUE    Unisphere-Egress-Statistics    disable 0
VALUE    Unisphere-Egress-Statistics    enable  1

#JUNOSe only
ATTRIBUTE Unisphere-Service-Category    ERX-VSA(14, integer) r
VALUE    Unisphere-Service-Category    UBR      1
VALUE    Unisphere-Service-Category    UBRPCR   2
VALUE    Unisphere-Service-Category    nrtVBR   3
VALUE    Unisphere-Service-Category    CBR      4
VALUE    Unisphere-Service-Category    rtVBR    5

#JUNOSe only
ATTRIBUTE Unisphere-pcr                  ERX-VSA(15, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-scr                  ERX-VSA(16, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-mbs                  ERX-VSA(17, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Init-CLI-Access-Level ERX-VSA(18, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Allow-All-VR-Access  ERX-VSA(19, integer) r
VALUE    Unisphere-Allow-All-VR-Access  disable 0
VALUE    Unisphere-Allow-All-VR-Access  enable  1

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-Access-Level ERX-VSA(20, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-VRouter-Name ERX-VSA(21, string) R

#JUNOSe only
ATTRIBUTE Unisphere-SA-Validate          ERX-VSA(22, integer) r
VALUE    Unisphere-SA-Validate          disable 0
VALUE    Unisphere-SA-Validate          enable  1

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Igmp-enable          ERX-VSA(23, integer) r
VALUE    Unisphere-Igmp-enable          disable 0
VALUE    Unisphere-Igmp-enable          enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description    ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name ERX-VSA(25, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name     ERX-VSA(26, string) r

#JUNOSe only
# ERX-VSA(27) reserved

#JUNOSe only
ATTRIBUTE Unisphere-PppoE-Url            ERX-VSA(28, string) r

#JUNOSe only
# ERX-VSA(29) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method      ERX-VSA(30, integer) r
VALUE    Unisphere-Nas-Port-Method      none 1
VALUE    Unisphere-Nas-Port-Method      cisco-avp 2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Bundle       ERX-VSA(31, string) r

#JUNOSe only
# ERX-VSA(32) reserved

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Max-Sessions  ERX-VSA(33, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Framed-Ip-Route-Tag  ERX-VSA(34, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Dialout-Number ERX-VSA(35, string) rt

```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Username          ERX-VSA(36, string) rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Password          ERX-VSA(37, string) rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Ppp-Protocol          ERX-VSA(38, integer) rt
VALUE    Unisphere-Ppp-Protocol          none      0
VALUE    Unisphere-Ppp-Protocol          pap        1
VALUE    Unisphere-Ppp-Protocol          chap       2
VALUE    Unisphere-Ppp-Protocol          pap-chap   3
VALUE    Unisphere-Ppp-Protocol          chap-pap   4
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Min-Bps        ERX-VSA(39, integer) rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Max-Bps        ERX-VSA(40, integer) rt
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Bearer-Type    ERX-VSA(41, integer) rt
VALUE    Unisphere-Tunnel-Bearer-Type    none      0
VALUE    Unisphere-Tunnel-Bearer-Type    analog   1
VALUE    Unisphere-Tunnel-Bearer-Type    digital  2
```

```
#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets     ERX-VSA(42, integer)
```

```
#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets    ERX-VSA(43, integer)
```

```
#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id   ERX-VSA(44, string)
```

```
#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router   ERX-VSA(45, string) r
```

```
#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface  ERX-VSA(46, string) r
```

```
#JUNOS and JUNOSe
```



```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number          ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only
ATTRIBUTE Unisphere-MLPPP-Bundle-Name        ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc           ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group             ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service       ERX-VSA(66, string) R

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1     ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1      ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1     ERX-TAGGED-INT-VSA(68, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1       ERX-TAGGED-INT-VSA(69, 1) rt
VALUE    Unisphere-Service-Stats-tag1       disabled          0
VALUE    Unisphere-Service-Stats-tag1       timestampOnly      1
VALUE    Unisphere-Service-Stats-tag1       timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1     ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2    ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2      ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2     ERX-TAGGED-INT-VSA(68, 2) rt

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2          ERX-TAGGED-INT-VSA(69, 2) rt
VALUE     Unisphere-Service-Stats-tag2          disabled          0
VALUE     Unisphere-Service-Stats-tag2          timestampOnly     1
VALUE     Unisphere-Service-Stats-tag2          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag2         ERX-TAGGED-INT-VSA(140, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag3        ERX-TAGGED-STRING-VSA(65, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag3         ERX-TAGGED-INT-VSA(67, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag3        ERX-TAGGED-INT-VSA(68, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag3          ERX-TAGGED-INT-VSA(69, 3) rt
VALUE     Unisphere-Service-Stats-tag3          disabled          0
VALUE     Unisphere-Service-Stats-tag3          timestampOnly     1
VALUE     Unisphere-Service-Stats-tag3          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3         ERX-TAGGED-INT-VSA(140, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag4        ERX-TAGGED-STRING-VSA(65, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag4         ERX-TAGGED-INT-VSA(67, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag4        ERX-TAGGED-INT-VSA(68, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag4          ERX-TAGGED-INT-VSA(69, 4) rt
VALUE     Unisphere-Service-Stats-tag4          disabled          0
VALUE     Unisphere-Service-Stats-tag4          timestampOnly     1
VALUE     Unisphere-Service-Stats-tag4          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4         ERX-TAGGED-INT-VSA(140, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag5        ERX-TAGGED-STRING-VSA(65, 5) rt

#JUNOSe only

```

```

ATTRIBUTE Unisphere-Service-Volume-tag5          ERX-TAGGED-INT-VSA(67, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag5          ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag5            ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5          disabled          0
VALUE      Unisphere-Service-Stats-tag5          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag5          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5          ERX-TAGGED-INT-VSA(140, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag6          ERX-TAGGED-STRING-VSA(65, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag6           ERX-TAGGED-INT-VSA(67, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag6           ERX-TAGGED-INT-VSA(68, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag6            ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6          disabled          0
VALUE      Unisphere-Service-Stats-tag6          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag6          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag6           ERX-TAGGED-INT-VSA(140, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag7          ERX-TAGGED-STRING-VSA(65, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag7           ERX-TAGGED-INT-VSA(67, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag7           ERX-TAGGED-INT-VSA(68, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag7            ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7          disabled          0
VALUE      Unisphere-Service-Stats-tag7          timestampOnly    1
VALUE      Unisphere-Service-Stats-tag7          timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7           ERX-TAGGED-INT-VSA(140, 7) rt

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag8      ERX-TAGGED-STRING-VSA(65, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag8        ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8       ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8         ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8       disabled          0
VALUE      Unisphere-Service-Stats-tag8       timestampOnly     1
VALUE      Unisphere-Service-Stats-tag8       timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8       ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                    ERX-VSA(70, integer) r
VALUE      Unisphere-DF-Bit                   dont-ignore-df-bit 0
VALUE      Unisphere-DF-Bit                   ignore-df-bit 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name          ERX-VSA(71, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Src-Name      ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name         ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name           ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name       ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name          ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version               ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version              v1 1

```



```

VALUE      Unisphere-MLD-Version          v2    2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version          ERX-VSA(78, integer) r
VALUE      Unisphere-IGMP-Version          v1    1
VALUE      Unisphere-IGMP-Version          v2    2
VALUE      Unisphere-IGMP-Version          v3    3

#JUNOSe only
ATTRIBUTE Unisphere-IP-Mcast-Adm-Bw-Limit ERX-VSA(79, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IPv6-Mcast-Adm-Bw-Limit ERX-VSA(80, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Access-Loop-Parameters ERX-VSA(81, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Parameters          ERX-VSA(82, string) R0

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Session          ERX-VSA(83, string) r

# Mobile-IP attributes

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Algorithm      ERX-VSA(84, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-SPI            ERX-VSA(85, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Key            ERX-VSA(86, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Replay          ERX-VSA(87, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Access-Control ERX-VSA(88, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Mobile-IP-Lifetime        ERX-VSA(89, integer) r

```

#JUNOSe only

ATTRIBUTE	Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE	Unisphere-L2TP-Resynch-Method	disable 0
VALUE	Unisphere-L2TP-Resynch-Method	failover 1
VALUE	Unisphere-L2TP-Resynch-Method	silent-failover 2
VALUE	Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3

#JUNOSe only

ATTRIBUTE	Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
-----------	---------------------------------	-----------------------

#JUNOSe only

ATTRIBUTE	Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
-----------	------------------------------	-----------------------

#JUNOSe only

ATTRIBUTE	Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
-----------	--------------------------------	-----------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE	Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE	Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE	Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE	Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE	Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE	Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE	Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE	Unisphere-Tunnel-Tx-Speed-Method	service-profile 7

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Query-Interval	ERX-VSA(95, integer) r
-----------	-------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Max-Resp-Time	ERX-VSA(96, integer) r
-----------	------------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-IGMP-Immediate-Leave	ERX-VSA(97, integer) r
VALUE	Unisphere-IGMP-Immediate-Leave	disabled 0
VALUE	Unisphere-IGMP-Immediate-Leave	enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Query-Interval	ERX-VSA(98, integer) r
-----------	------------------------------	------------------------

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Max-Resp-Time	ERX-VSA(99, integer) r
-----------	-----------------------------	------------------------

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-MLD-Immediate-Leave	ERX-VSA(100, integer)	r
VALUE	Unisphere-MLD-Immediate-Leave	disabled	0
VALUE	Unisphere-MLD-Immediate-Leave	enabled	1

#JUNOSe only

ATTRIBUTE	Unisphere-IP-Block-Multicast	ERX-VSA(101, integer)	r
VALUE	Unisphere-IP-Block-Multicast	disabled	0
VALUE	Unisphere-IP-Block-Multicast	enabled	1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-Explicit-Tracking	ERX-VSA(102, integer)	r
VALUE	Unisphere-IGMP-Explicit-Tracking	disabled	0
VALUE	Unisphere-IGMP-Explicit-Tracking	enabled	1

#JUNOSe only

ATTRIBUTE	Unisphere-IGMP-No-Tracking-V2-Grps	ERX-VSA(103, integer)	r
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	disabled	0
VALUE	Unisphere-IGMP-No-Tracking-V2-Grps	enabled	1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-Explicit-Tracking	ERX-VSA(104, integer)	r
VALUE	Unisphere-MLD-Explicit-Tracking	disabled	0
VALUE	Unisphere-MLD-Explicit-Tracking	enabled	1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer)	r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled	0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled	1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string)	r
-----------	-------------------------------	----------------------	---

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string)	r
-----------	------------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string)	R
-----------	-------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr)	R
-----------	-------------------------------	----------------------	---

#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Up	ERX-VSA(119, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Data-Rate-Dn	ERX-VSA(120, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up	ERX-VSA(121, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn	ERX-VSA(122, integer) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up	ERX-VSA(123, integer) c
#JUNOS and JUNOSe	

```

ATTRIBUTE Unisphere-Act-Interlv-Delay-Up          ERX-VSA(124, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn          ERX-VSA(125, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn          ERX-VSA(126, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Line-State                ERX-VSA(127, integer) c
VALUE      Unisphere-DSL-Line-State              SHOWTIME 1
VALUE      Unisphere-DSL-Line-State              IDLE     2
VALUE      Unisphere-DSL-Line-State              SILENT    3

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Type                      ERX-VSA(128, integer) c
VALUE      Unisphere-DSL-Type                    ADSL1      1
VALUE      Unisphere-DSL-Type                    ADSL2      2
VALUE      Unisphere-DSL-Type                    ADSL2PLUS  3
VALUE      Unisphere-DSL-Type                    VDSL1      4
VALUE      Unisphere-DSL-Type                    VDSL2      5
VALUE      Unisphere-DSL-Type                    SDSL       6
VALUE      Unisphere-DSL-Type                    UNKNOWN    7
VALUE      Unisphere-DSL-Type                    G.Fast     8
VALUE      Unisphere-DSL-Type                    VDSL2-AnnexQ 9
VALUE      Unisphere-DSL-Type                    SDSL-Bonded 10
VALUE      Unisphere-DSL-Type                    VDSL2-Bonded 11
VALUE      Unisphere-DSL-Type                    G.Fast-Bonded 12
VALUE      Unisphere-DSL-Type                    VDSL2-AnnexQ-Bonded 13

#JUNOSe only
ATTRIBUTE Unisphere-Ipv6-NdRa-Prefix              ERX-VSA(129, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Qos-Set-Name                  ERX-VSA(130, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DownStream-Calc-Rate          ERX-VSA(141, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-UpStream-Calc-Rate            ERX-VSA(142, integer) c

```

```

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-Max-Clients-Per-Interface          ERX-VSA(143, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-PPP-Ingress-Only              ERX-VSA(144, integer) r
VALUE      Unisphere-PPP-Ingress-Only            disabled  0
VALUE      Unisphere-PPP-Ingress-Only            enabled   1

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-Cos-Scheduler-Pmt-Type             ERX-VSA(146, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Backup-Address-Pool           ERX-VSA(147, string) r

#JUNOSe only
ATTRIBUTE Unisphere-ICR-Partition-Id              ERX-VSA(150, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Octets         ERX-VSA(151, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Octets        ERX-VSA(152, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Packets        ERX-VSA(153, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Packets       ERX-VSA(154, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Gigawords      ERX-VSA(155, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Gigawords     ERX-VSA(156, integer)

#JUNOSe only
ATTRIBUTE Jnpr-IPv6-Ndra-Pool-Name                 ERX-VSA(157, string) r

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-PppoE-Padn                          ERX-VSA(158, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Dhcp-Option-82                 ERX-VSA(159, string) c

```

```

#JUNOS only
ATTRIBUTE Jnpr-Vlan-Map-Id ERX-VSA(160, integer)

#JUNOS only
ATTRIBUTE Jnpr-IPv6-Delegated-Pool-Name ERX-VSA(161, string) r

#JUNOS only
ATTRIBUTE Jnpr-Tx-Connect-Speed ERX-VSA(162, integer) r

#JUNOS only
ATTRIBUTE Jnpr-Rx-Connect-Speed ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IPv4-release-control ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP
# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type ERX-VSA(173, integer) rt
VALUE Unisphere-Service-Activate-Type dynamic-profile 1

```

VALUE	Unisphere-Service-Activate-Type	op-script 2
#JUNOS only		
ATTRIBUTE	Unisphere-Client-Profile-Name	ERX-VSA(174, string) r
#JUNOS only - MobileNext Gateway only		
# ERX-VSA(175) reserved for Redirect-Gw-Addr		
#JUNOS only - MobileNext Gateway only		
# ERX-VSA(176) reserved for Apn-Name		
#JUNOS only		
ATTRIBUTE	Unisphere-Cos-Shaping-Rate	ERX-VSA(177, string)
#JUNOS only		
ATTRIBUTE	Unisphere-Action-Reason	ERX-VSA(178, string)
#JUNOS only		
ATTRIBUTE	Unisphere-Service-Volume-Gigawords	ERX-VSA(179, integer) rt
#JUNOS only		
ATTRIBUTE	Unisphere-Update-Service	ERX-VSA(180, string) rt
#JUNOS only		
# ERX-VSA(181) reserved for 15.1		
#JUNOS only		
ATTRIBUTE	Unisphere-Acc-Loop-Remote-Id	ERX-VSA(182, string) c
#JUNOS only		
ATTRIBUTE	Unisphere-Acc-Loop-Encap	ERX-VSA(183, hexadecimal) c
#JUNOS only		
ATTRIBUTE	Unisphere-Inner-Vlan-Map-Id	ERX-VSA(184, integer) c
#JUNOS only		
ATTRIBUTE	Unisphere-Core-Facing-Interface	ERX-VSA(185, string) c
#JUNOS only		
ATTRIBUTE	Unisphere-Pcp-Port-Map	ERX-VSA(186, string) R
#JUNOS only		
ATTRIBUTE	Unisphere-vCPE-Lan-Extension	ERX-VSA(187, string) r

#JUNOS only

ATTRIBUTE Unisphere-vCPE-IPv4-Offload ERX-VSA(188, string) r

#JUNOS only

ATTRIBUTE Jnpr-Input-Interface-Filter ERX-VSA(191, string) r

ATTRIBUTE Jnpr-Output-Interface-Filter ERX-VSA(192, string) r

#JUNOS only

ATTRIBUTE ERX-Bulk-CoA-Transaction-Id ERX-VSA(194, integer) r

ATTRIBUTE ERX-Bulk-CoA-Identifier ERX-VSA(195, integer) r

#JUNOS only

ATTRIBUTE Unisphere-IPv4-Input-Service-Set ERX-VSA(196, string) r

ATTRIBUTE Unisphere-IPv4-Output-Service-Set ERX-VSA(197, string) r

ATTRIBUTE Unisphere-IPv4-Input-Service-Filter ERX-VSA(198, string) r

ATTRIBUTE Unisphere-IPv4-Output-Service-Filter ERX-VSA(199, string) r

ATTRIBUTE Unisphere-IPv6-Input-Service-Set ERX-VSA(200, string) r

ATTRIBUTE Unisphere-IPv6-Output-Service-Set ERX-VSA(201, string) r

ATTRIBUTE Unisphere-IPv6-Input-Service-Filter ERX-VSA(202, string) r

ATTRIBUTE Unisphere-IPv6-Output-Service-Filter ERX-VSA(203, string) r

ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name ERX-VSA(204, string) r

ATTRIBUTE Unisphere-Adv-Pcef-Rule-Name ERX-VSA(205, string) R

#JUNOS only

ATTRIBUTE ERX-DHCPv6-Options ERX-VSA(207, hexadecimal) CRO

ATTRIBUTE ERX-DHCP-Header ERX-VSA(208, hexadecimal) c

ATTRIBUTE ERX-DHCPv6-Header ERX-VSA(209, hexadecimal) c

ATTRIBUTE ERX-Acct-Request-Reason ERX-VSA(210, integer) c

VALUE ERX-Acct-Request-Reason Acct-Start-Ack 1

VALUE ERX-Acct-Request-Reason Periodic-Interval-Interim 2

VALUE ERX-Acct-Request-Reason IP-Active 4

VALUE ERX-Acct-Request-Reason IP-Inactive 8

VALUE ERX-Acct-Request-Reason IPv6-Active 16

VALUE ERX-Acct-Request-Reason IPv6-Inactive 32

VALUE ERX-Acct-Request-Reason Session-Active 64

VALUE ERX-Acct-Request-Reason Session-Inactive 128

VALUE ERX-Acct-Request-Reason Line-Speed-Change 256

VALUE ERX-Acct-Request-Reason Address-Assignment-Change 512

VALUE ERX-Acct-Request-Reason CoA-Processing-Complete 1024

ATTRIBUTE ERX-Inner-Vlan-Tag-Protocol-Id ERX-VSA(211, string) c

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Routing-Services	ERX-VSA(212, integer) r
VALUE	Unisphere-Routing-Services	disable 0
VALUE	Unisphere-Routing-Services	enable 1

#JUNOS and JUNOSe

tagged tunnel attributes

ATTRIBUTE	Unisphere-Tunnel-VR-tag1	ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag1	ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag1	ERX-TAGGED-INT-VSA(33, 1) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag2	ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag2	ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag2	ERX-TAGGED-INT-VSA(33, 2) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag3	ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag3	ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag3	ERX-TAGGED-INT-VSA(33, 3) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag4	ERX-TAGGED-STRING-VSA(8, 4) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt

ATTRIBUTE	Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE	Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE	Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt

ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29	ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29	ERX-TAGGED-INT-VSA(33, 29) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag30	ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30	ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30	ERX-TAGGED-INT-VSA(33, 30) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag31	ERX-TAGGED-STRING-VSA(8, 31) rt

```

ATTRIBUTE Unisphere-Tunnel-Password-tag31      ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31      ERX-TAGGED-INT-VSA(33, 31) rt

#JUNOS only
##### tagged service activation type attributes

ATTRIBUTE Unisphere-Service-Act-Type-tag1      ERX-TAGGED-INT-VSA(173, 1) rt
VALUE      Unisphere-Service-Act-Type-tag1      dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag1      op-script 2

ATTRIBUTE Unisphere-Service-Act-Type-tag2      ERX-TAGGED-INT-VSA(173, 2) rt
VALUE      Unisphere-Service-Act-Type-tag2      dynamic-profile 1
VALUE      Unisphere-Service-Act-Type-tag2      op-script 2

#####
# unisphere.dct - Unisphere ERX dictionary
#####

```

Junos OS Release 18.4 Subscriber Management RADIUS Dictionary [DCT]

IN THIS SECTION

- [Junos OS Release 18.4 Subscriber Management RADIUS Dictionary \[DCT\] | 990](#)

Junos OS Release 18.4 Subscriber Management RADIUS Dictionary [DCT]

```

#####
# unisphere.dct - Unisphere MX Family dictionary to support JUNOS v18.4
# (Note: JUNOSe attribute breakdown by JUNOSe releases not reflected in this document.
#       JUNOSe information is based on complete JUNOSe attribute listing as of 03/18/16).
#
#
#

```

```

# (See README.DCT for more details on the format of this file)
#
#   please note that tunnel attributes (8, 9, 33, 35, 39 - 41)
#       and service attributes (65 - 69) may be tagged
#
#####
#
# Use the Radius specification attributes
#
@radius.dct

#
# Define additional Unisphere ERX Family Attributes
#

# for untagged attributes: t = attr type value (attr id); s = data type (integer, string, etc)
MACRO ERX-VSA(t,s) 26 [vid=4874 type1=%t% len1=+2 data=%s%]

# for tagged strings t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-STRING-VSA(t,f) 26 [vid=4874 type1=%t% len1=+3 tag=%f% data=string]

# for integers t = attr type value (attr id); f = tag value
MACRO ERX-TAGGED-INT-VSA(t,f) 26 [vid=4874 type1=%t% len1=+2 tag=%f% data=integer]

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Virtual-Router          ERX-VSA(1, string)  rt

#JUNOSe only
ATTRIBUTE Unisphere-Local-Address-Pool      ERX-VSA(2, string)  r

#JUNOSe only
ATTRIBUTE Unisphere-Local-Interface         ERX-VSA(3, string)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Primary-Dns             ERX-VSA(4, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Secondary-Dns           ERX-VSA(5, ipaddr)  r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Primary-Wins            ERX-VSA(6, ipaddr)  r

#JUNOS and JUNOSe

```

ATTRIBUTE Unisphere-Secondary-Wins	ERX-VSA(7, ipaddr) r
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Tunnel-Virtual-Router	ERX-VSA(8, string) rt
#JUNOS and JUNOSe	
# untagged	
ATTRIBUTE Unisphere-Tunnel-Password	ERX-VSA(9, string) rt
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Ingress-Policy-Name	ERX-VSA(10, string) r
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Egress-Policy-Name	ERX-VSA(11, string) r
#JUNOSe only	
ATTRIBUTE Unisphere-Ingress-Statistics	ERX-VSA(12, integer) r
VALUE Unisphere-Ingress-Statistics	disable 0
VALUE Unisphere-Ingress-Statistics	enable 1
#JUNOSe only	
ATTRIBUTE Unisphere-Egress-Statistics	ERX-VSA(13, integer) r
VALUE Unisphere-Egress-Statistics	disable 0
VALUE Unisphere-Egress-Statistics	enable 1
#JUNOSe only	
ATTRIBUTE Unisphere-Service-Category	ERX-VSA(14, integer) r
VALUE Unisphere-Service-Category	UBR 1
VALUE Unisphere-Service-Category	UBRPCR 2
VALUE Unisphere-Service-Category	nrtVBR 3
VALUE Unisphere-Service-Category	CBR 4
VALUE Unisphere-Service-Category	rtVBR 5
#JUNOSe only	
ATTRIBUTE Unisphere-pcr	ERX-VSA(15, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-scr	ERX-VSA(16, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-mbs	ERX-VSA(17, integer) r
#JUNOSe only	

```

ATTRIBUTE Unisphere-Init-CLI-Access-Level      ERX-VSA(18, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Allow-All-VR-Access        ERX-VSA(19, integer) r
VALUE      Unisphere-Allow-All-VR-Access      disable 0
VALUE      Unisphere-Allow-All-VR-Access      enable 1

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-Access-Level        ERX-VSA(20, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Alt-CLI-VRouter-Name        ERX-VSA(21, string) R

#JUNOSe only
ATTRIBUTE Unisphere-SA-Validate                 ERX-VSA(22, integer) r
VALUE      Unisphere-SA-Validate               disable 0
VALUE      Unisphere-SA-Validate               enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Igmp-enable                 ERX-VSA(23, integer) r
VALUE      Unisphere-Igmp-enable               disable 0
VALUE      Unisphere-Igmp-enable               enable 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Pppoe-Description           ERX-VSA(24, string)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Redirect-VRouter-Name       ERX-VSA(25, string) rt

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Profile-Name           ERX-VSA(26, string) r

#JUNOSe only
# ERX-VSA(27) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Pppoe-Ur1                  ERX-VSA(28, string) r

#JUNOSe only
# ERX-VSA(29) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Nas-Port-Method             ERX-VSA(30, integer) r

```


VALUE	Unisphere-Nas-Port-Method	none 1
VALUE	Unisphere-Nas-Port-Method	cisco-avp 2
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Service-Bundle	ERX-VSA(31, string) r
#JUNOSe only		
# ERX-VSA(32) reserved		
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Tunnel-Max-Sessions	ERX-VSA(33, integer) rt
#JUNOS and JUNOSe		
ATTRIBUTE	Unisphere-Framed-Ip-Route-Tag	ERX-VSA(34, integer) r
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Dialout-Number	ERX-VSA(35, string) rt
#JUNOSe only		
ATTRIBUTE	Unisphere-Ppp-Username	ERX-VSA(36, string) rt
#JUNOSe only		
ATTRIBUTE	Unisphere-Ppp-Password	ERX-VSA(37, string) rt
#JUNOSe only		
ATTRIBUTE	Unisphere-Ppp-Protocol	ERX-VSA(38, integer) rt
VALUE	Unisphere-Ppp-Protocol	none 0
VALUE	Unisphere-Ppp-Protocol	pap 1
VALUE	Unisphere-Ppp-Protocol	chap 2
VALUE	Unisphere-Ppp-Protocol	pap-chap 3
VALUE	Unisphere-Ppp-Protocol	chap-pap 4
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Min-Bps	ERX-VSA(39, integer) rt
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Max-Bps	ERX-VSA(40, integer) rt
#JUNOSe only		
ATTRIBUTE	Unisphere-Tunnel-Bearer-Type	ERX-VSA(41, integer) rt
VALUE	Unisphere-Tunnel-Bearer-Type	none 0
VALUE	Unisphere-Tunnel-Bearer-Type	analog 1
VALUE	Unisphere-Tunnel-Bearer-Type	digital 2

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Input-Gigapackets          ERX-VSA(42, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Output-Gigapackets         ERX-VSA(43, integer)

#JUNOSe only
ATTRIBUTE Unisphere-Tunnel-Interface-Id        ERX-VSA(44, string)

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Virtual-Router        ERX-VSA(45, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IpV6-Local-Interface       ERX-VSA(46, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Primary-Dns           ERX-VSA(47, hexadecimal) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Secondary-Dns         ERX-VSA(48, hexadecimal) r

#JUNOSe only
# ERX-VSA(49) reserved

#JUNOSe only
# ERX-VSA(50) reserved

#JUNOSe only
ATTRIBUTE Unisphere-Disconnect-Cause           ERX-VSA(51, hexadecimal)

#JUNOSe only
ATTRIBUTE Unisphere-Radius-Client-Address      ERX-VSA(52, ipaddr)

#JUNOSe only
ATTRIBUTE Unisphere-Service-Description        ERX-VSA(53, string) c

#JUNOSe only
ATTRIBUTE Unisphere-L2tp-Recv-Window-Size     ERX-VSA(54, integer) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Options               ERX-VSA(55, hexadecimal) cRO

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Mac-Addr          ERX-VSA(56, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Dhcp-Gi-Address        ERX-VSA(57, ipaddr) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-LI-Action              ERX-VSA(58, integer) rt salt-encrypt
VALUE    Unisphere-LI-Action              off    0
VALUE    Unisphere-LI-Action              on     1
VALUE    Unisphere-LI-Action              noop   2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Dev-Handle         ERX-VSA(59, hexadecimal) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Ip-Address         ERX-VSA(60, ipaddr) rt salt-encrypt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Med-Port-Number        ERX-VSA(61, integer) rt salt-encrypt

#JUNOSe only
ATTRIBUTE Unisphere-MLPPP-Bundle-Name      ERX-VSA(62, string) c

#JUNOS only
ATTRIBUTE Unisphere-Interface-Desc         ERX-VSA(63, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Tunnel-Group           ERX-VSA(64, string) r

# please note that these service attributes may be tagged

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Deactivate-Service     ERX-VSA(66, string) R

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag1  ERX-TAGGED-STRING-VSA(65, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag1    ERX-TAGGED-INT-VSA(67, 1) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag1   ERX-TAGGED-INT-VSA(68, 1) rt

```

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag1      ERX-TAGGED-INT-VSA(69, 1) rt
VALUE     Unisphere-Service-Stats-tag1      disabled          0
VALUE     Unisphere-Service-Stats-tag1      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag1      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag1     ERX-TAGGED-INT-VSA(140, 1) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag2    ERX-TAGGED-STRING-VSA(65, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag2     ERX-TAGGED-INT-VSA(67, 2) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag2     ERX-TAGGED-INT-VSA(68, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag2      ERX-TAGGED-INT-VSA(69, 2) rt
VALUE     Unisphere-Service-Stats-tag2      disabled          0
VALUE     Unisphere-Service-Stats-tag2      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag2      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag2     ERX-TAGGED-INT-VSA(140, 2) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag3    ERX-TAGGED-STRING-VSA(65, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag3     ERX-TAGGED-INT-VSA(67, 3) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag3     ERX-TAGGED-INT-VSA(68, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag3      ERX-TAGGED-INT-VSA(69, 3) rt
VALUE     Unisphere-Service-Stats-tag3      disabled          0
VALUE     Unisphere-Service-Stats-tag3      timestampOnly     1
VALUE     Unisphere-Service-Stats-tag3      timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag3     ERX-TAGGED-INT-VSA(140, 3) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag4    ERX-TAGGED-STRING-VSA(65, 4) rt

```

```

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag4          ERX-TAGGED-INT-VSA(67, 4) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag4         ERX-TAGGED-INT-VSA(68, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag4           ERX-TAGGED-INT-VSA(69, 4) rt
VALUE      Unisphere-Service-Stats-tag4         disabled          0
VALUE      Unisphere-Service-Stats-tag4         timestampOnly    1
VALUE      Unisphere-Service-Stats-tag4         timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag4         ERX-TAGGED-INT-VSA(140, 4) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag5        ERX-TAGGED-STRING-VSA(65, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag5          ERX-TAGGED-INT-VSA(67, 5) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag5         ERX-TAGGED-INT-VSA(68, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag5           ERX-TAGGED-INT-VSA(69, 5) rt
VALUE      Unisphere-Service-Stats-tag5         disabled          0
VALUE      Unisphere-Service-Stats-tag5         timestampOnly    1
VALUE      Unisphere-Service-Stats-tag5         timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag5         ERX-TAGGED-INT-VSA(140, 5) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag6        ERX-TAGGED-STRING-VSA(65, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag6          ERX-TAGGED-INT-VSA(67, 6) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag6         ERX-TAGGED-INT-VSA(68, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag6           ERX-TAGGED-INT-VSA(69, 6) rt
VALUE      Unisphere-Service-Stats-tag6         disabled          0
VALUE      Unisphere-Service-Stats-tag6         timestampOnly    1
VALUE      Unisphere-Service-Stats-tag6         timestampAndVolume 2

```

```

ATTRIBUTE Unisphere-Service-AcctInt-tag6          ERX-TAGGED-INT-VSA(140, 6) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag7          ERX-TAGGED-STRING-VSA(65, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag7            ERX-TAGGED-INT-VSA(67, 7) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag7           ERX-TAGGED-INT-VSA(68, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag7             ERX-TAGGED-INT-VSA(69, 7) rt
VALUE      Unisphere-Service-Stats-tag7           disabled          0
VALUE      Unisphere-Service-Stats-tag7           timestampOnly     1
VALUE      Unisphere-Service-Stats-tag7           timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag7           ERX-TAGGED-INT-VSA(140, 7) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Activate-Service-tag8          ERX-TAGGED-STRING-VSA(65, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Volume-tag8            ERX-TAGGED-INT-VSA(67, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-Service-Timeout-tag8           ERX-TAGGED-INT-VSA(68, 8) rt

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Stats-tag8             ERX-TAGGED-INT-VSA(69, 8) rt
VALUE      Unisphere-Service-Stats-tag8           disabled          0
VALUE      Unisphere-Service-Stats-tag8           timestampOnly     1
VALUE      Unisphere-Service-Stats-tag8           timestampAndVolume 2
ATTRIBUTE Unisphere-Service-AcctInt-tag8           ERX-TAGGED-INT-VSA(140, 8) rt

#JUNOSe only
ATTRIBUTE Unisphere-DF-Bit                          ERX-VSA(70, integer) r
VALUE      Unisphere-DF-Bit                        dont-ignore-df-bit 0
VALUE      Unisphere-DF-Bit                        ignore-df-bit 1

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Access-Name               ERX-VSA(71, string) r

#JUNOS and JUNOSe

```

```

ATTRIBUTE Unisphere-IGMP-Access-Src-Name          ERX-VSA(72, string) r

#JUNOSe only
ATTRIBUTE Unisphere-IGMP-OIF-Map-Name             ERX-VSA(73, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Name               ERX-VSA(74, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Access-Src-Name           ERX-VSA(75, string) r

#JUNOSe only
ATTRIBUTE Unisphere-MLD-OIF-Map-Name              ERX-VSA(76, string) r

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-MLD-Version                   ERX-VSA(77, integer) r
VALUE      Unisphere-MLD-Version                  v1      1
VALUE      Unisphere-MLD-Version                  v2      2

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-IGMP-Version                  ERX-VSA(78, integer) r
VALUE      Unisphere-IGMP-Version                  v1      1
VALUE      Unisphere-IGMP-Version                  v2      2
VALUE      Unisphere-IGMP-Version                  v3      3

#JUNOSe only
ATTRIBUTE Unisphere-IP-Mcast-Adm-Bw-Limit         ERX-VSA(79, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-IPv6-Mcast-Adm-Bw-Limit       ERX-VSA(80, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-L2c-Access-Loop-Parameters    ERX-VSA(81, string) r

#JUNOSe only
ATTRIBUTE Unisphere-Qos-Parameters                ERX-VSA(82, string) R0

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Service-Session               ERX-VSA(83, string) r

# Mobile-IP attributes

#JUNOSe only

```

ATTRIBUTE Unisphere-Mobile-IP-Algorithm	ERX-VSA(84, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-Mobile-IP-SPI	ERX-VSA(85, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-Mobile-IP-Key	ERX-VSA(86, string) r
#JUNOSe only	
ATTRIBUTE Unisphere-Mobile-IP-Replay	ERX-VSA(87, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-Mobile-IP-Access-Control	ERX-VSA(88, string) r
#JUNOSe only	
ATTRIBUTE Unisphere-Mobile-IP-Lifetime	ERX-VSA(89, integer) r
#JUNOSe only	
ATTRIBUTE Unisphere-L2TP-Resynch-Method	ERX-VSA(90, integer) r
VALUE Unisphere-L2TP-Resynch-Method	disable 0
VALUE Unisphere-L2TP-Resynch-Method	failover 1
VALUE Unisphere-L2TP-Resynch-Method	silent-failover 2
VALUE Unisphere-L2TP-Resynch-Method	failover-with-silent-backup 3
#JUNOSe only	
ATTRIBUTE Unisphere-Tunnel-Switch-Profile	ERX-VSA(91, string) r
#JUNOSe only	
ATTRIBUTE Unisphere-L2c-Up-Stream-Data	ERX-VSA(92, string) c
#JUNOSe only	
ATTRIBUTE Unisphere-L2c-Down-Stream-Data	ERX-VSA(93, string) c
#JUNOS and JUNOSe	
ATTRIBUTE Unisphere-Tunnel-Tx-Speed-Method	ERX-VSA(94, integer) rt
VALUE Unisphere-Tunnel-Tx-Speed-Method	none 0
VALUE Unisphere-Tunnel-Tx-Speed-Method	static-layer2 1
VALUE Unisphere-Tunnel-Tx-Speed-Method	dynamic-layer2 2
VALUE Unisphere-Tunnel-Tx-Speed-Method	qos 3
VALUE Unisphere-Tunnel-Tx-Speed-Method	actual 4
VALUE Unisphere-Tunnel-Tx-Speed-Method	ancp 5
VALUE Unisphere-Tunnel-Tx-Speed-Method	pppoe-ia-tags 6
VALUE Unisphere-Tunnel-Tx-Speed-Method	service-profile 7

#JUNOSe only

ATTRIBUTE Unisphere-IGMP-Query-Interval ERX-VSA(95, integer) r

#JUNOSe only

ATTRIBUTE Unisphere-IGMP-Max-Resp-Time ERX-VSA(96, integer) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-IGMP-Immediate-Leave ERX-VSA(97, integer) r

VALUE Unisphere-IGMP-Immediate-Leave disabled 0

VALUE Unisphere-IGMP-Immediate-Leave enabled 1

#JUNOSe only

ATTRIBUTE Unisphere-MLD-Query-Interval ERX-VSA(98, integer) r

#JUNOSe only

ATTRIBUTE Unisphere-MLD-Max-Resp-Time ERX-VSA(99, integer) r

#JUNOS and JUNOSe

ATTRIBUTE Unisphere-MLD-Immediate-Leave ERX-VSA(100, integer) r

VALUE Unisphere-MLD-Immediate-Leave disabled 0

VALUE Unisphere-MLD-Immediate-Leave enabled 1

#JUNOSe only

ATTRIBUTE Unisphere-IP-Block-Multicast ERX-VSA(101, integer) r

VALUE Unisphere-IP-Block-Multicast disabled 0

VALUE Unisphere-IP-Block-Multicast enabled 1

#JUNOSe only

ATTRIBUTE Unisphere-IGMP-Explicit-Tracking ERX-VSA(102, integer) r

VALUE Unisphere-IGMP-Explicit-Tracking disabled 0

VALUE Unisphere-IGMP-Explicit-Tracking enabled 1

#JUNOSe only

ATTRIBUTE Unisphere-IGMP-No-Tracking-V2-Grps ERX-VSA(103, integer) r

VALUE Unisphere-IGMP-No-Tracking-V2-Grps disabled 0

VALUE Unisphere-IGMP-No-Tracking-V2-Grps enabled 1

#JUNOSe only

ATTRIBUTE Unisphere-MLD-Explicit-Tracking ERX-VSA(104, integer) r

VALUE Unisphere-MLD-Explicit-Tracking disabled 0

VALUE Unisphere-MLD-Explicit-Tracking enabled 1

#JUNOSe only

ATTRIBUTE	Unisphere-MLD-No-Tracking-V1-Grps	ERX-VSA(105, integer)	r
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	disabled	0
VALUE	Unisphere-MLD-No-Tracking-V1-Grps	enabled	1

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Ingress-Policy-Name	ERX-VSA(106, string)	r
-----------	-------------------------------	----------------------	---

#JUNOS only

ATTRIBUTE	Jnpr-IPv6-Egress-Policy-Name	ERX-VSA(107, string)	r
-----------	------------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-CoS-Parameter-Type	ERX-VSA(108, string)	R
-----------	-------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Jnpr-DHCP-Guided-Relay-Server	ERX-VSA(109, ipaddr)	R
-----------	-------------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Loop-Cir-Id	ERX-VSA(110, string)	c
-----------	---------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Bin	ERX-VSA(111, hexadecimal)	c
-----------	-------------------------------	---------------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Acc-Aggr-Cir-Id-Asc	ERX-VSA(112, string)	c
-----------	-------------------------------	----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Act-Data-Rate-Up	ERX-VSA(113, integer)	c
-----------	----------------------------	-----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Act-Data-Rate-Dn	ERX-VSA(114, integer)	c
-----------	----------------------------	-----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Min-Data-Rate-Up	ERX-VSA(115, integer)	c
-----------	----------------------------	-----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Min-Data-Rate-Dn	ERX-VSA(116, integer)	c
-----------	----------------------------	-----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Att-Data-Rate-Up	ERX-VSA(117, integer)	c
-----------	----------------------------	-----------------------	---

#JUNOS and JUNOSe

ATTRIBUTE	Unisphere-Att-Data-Rate-Dn	ERX-VSA(118, integer)	c
-----------	----------------------------	-----------------------	---

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Up          ERX-VSA(119, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Data-Rate-Dn          ERX-VSA(120, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Up        ERX-VSA(121, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Min-LP-Data-Rate-Dn        ERX-VSA(122, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Up       ERX-VSA(123, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Up       ERX-VSA(124, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Max-Interlv-Delay-Dn       ERX-VSA(125, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Act-Interlv-Delay-Dn       ERX-VSA(126, integer) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Line-State             ERX-VSA(127, integer) c
VALUE      Unisphere-DSL-Line-State           SHOWTIME 1
VALUE      Unisphere-DSL-Line-State           IDLE      2
VALUE      Unisphere-DSL-Line-State           SILENT     3

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-DSL-Type                   ERX-VSA(128, integer) c
VALUE      Unisphere-DSL-Type                  ADSL1             1
VALUE      Unisphere-DSL-Type                  ADSL2             2
VALUE      Unisphere-DSL-Type                  ADSL2PLUS         3
VALUE      Unisphere-DSL-Type                  VDSL1             4
VALUE      Unisphere-DSL-Type                  VDSL2             5
VALUE      Unisphere-DSL-Type                  SDSL              6
VALUE      Unisphere-DSL-Type                  UNKNOWN          7
VALUE      Unisphere-DSL-Type                  G.Fast            8
VALUE      Unisphere-DSL-Type                  VDSL2-AnnexQ      9
VALUE      Unisphere-DSL-Type                  SDSL-Bonded       10

```

VALUE	Unisphere-DSL-Type	VDSL2-Bonded	11
VALUE	Unisphere-DSL-Type	G.Fast-Bonded	12
VALUE	Unisphere-DSL-Type	VDSL2-AnnexQ-Bonded	13
#JUNOSe only			
ATTRIBUTE	Unisphere-Ipv6-NdRa-Prefix	ERX-VSA(129, hexadecimal)	r
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Qos-Set-Name	ERX-VSA(130, string)	r
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-DownStream-Calc-Rate	ERX-VSA(141, integer)	c
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-UpStream-Calc-Rate	ERX-VSA(142, integer)	c
#JUNOS and JUNOSe			
ATTRIBUTE	Jnpr-Max-Clients-Per-Interface	ERX-VSA(143, integer)	r
#JUNOSe only			
ATTRIBUTE	Unisphere-PPP-Ingress-Only	ERX-VSA(144, integer)	r
VALUE	Unisphere-PPP-Ingress-Only	disabled	0
VALUE	Unisphere-PPP-Ingress-Only	enabled	1
#JUNOS and JUNOSe			
ATTRIBUTE	Jnpr-Cos-Scheduler-Pmt-Type	ERX-VSA(146, string)	R
#JUNOSe only			
ATTRIBUTE	Unisphere-Backup-Address-Pool	ERX-VSA(147, string)	r
#JUNOSe only			
ATTRIBUTE	Unisphere-ICR-Partition-Id	ERX-VSA(150, string)	
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Ipv6-Acct-Input-Octets	ERX-VSA(151, integer)	
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Ipv6-Acct-Output-Octets	ERX-VSA(152, integer)	
#JUNOS and JUNOSe			
ATTRIBUTE	Unisphere-Ipv6-Acct-Input-Packets	ERX-VSA(153, integer)	

```

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Packets      ERX-VSA(154, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Input-Gigawords     ERX-VSA(155, integer)

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Ipv6-Acct-Output-Gigawords    ERX-VSA(156, integer)

#JUNOSe only
ATTRIBUTE Jnpr-IPv6-Ndra-Pool-Name                ERX-VSA(157, string) r

#JUNOS and JUNOSe
ATTRIBUTE Jnpr-PppoE-Padn                         ERX-VSA(158, string) R

#JUNOSe only
ATTRIBUTE Unisphere-Dhcp-Option-82               ERX-VSA(159, string) c

#JUNOS only
ATTRIBUTE Jnpr-Vlan-Map-Id                       ERX-VSA(160, integer)

#JUNOS only
ATTRIBUTE Jnpr-IPv6-Delegated-Pool-Name          ERX-VSA(161, string) r

#JUNOS only
ATTRIBUTE Jnpr-Tx-Connect-Speed                  ERX-VSA(162, integer) r

#JUNOS only
ATTRIBUTE Jnpr-Rx-Connect-Speed                  ERX-VSA(163, integer) r

#JUNOSe only
ATTRIBUTE Unisphere-Ipv4-release-control          ERX-VSA(164, string) r

#JUNOSe only - Pcp-Server-Name
ATTRIBUTE Pcp-Server-Name                        ERX-VSA(165, string) r

#JUNOS only - MSR JNPR-CGN-Event-Type
# ERX-VSA(166) reserved

#JUNOS only - MSR JNPR-CGN-Private-IP
# ERX-VSA(167) reserved

#JUNOS only - MSR JNPR-CGN-Public-IP

```

```

# ERX-VSA(168) reserved

#JUNOS only - MSR JNPR-CGN-Start-Port
# ERX-VSA(169) reserved

#JUNOS only - MSR JNPR-CGN-End-Port
# ERX-VSA(170) reserved

#JUNOS only - MSR JNPR-CGN-NAT-Domain
# ERX-VSA(171) reserved

#JUNOS only - MSR JNPR-CGN-Miscellaneous
# ERX-VSA(172) reserved

#JUNOS only
ATTRIBUTE Unisphere-Service-Activate-Type      ERX-VSA(173, integer) rt
VALUE     Unisphere-Service-Activate-Type      dynamic-profile 1
VALUE     Unisphere-Service-Activate-Type      op-script 2

#JUNOS only
ATTRIBUTE Unisphere-Client-Profile-Name         ERX-VSA(174, string) r

#JUNOS only - MobileNext Gateway only
# ERX-VSA(175) reserved for Redirect-Gw-Addr

#JUNOS only - MobileNext Gateway only
# ERX-VSA(176) reserved for Apn-Name

#JUNOS only
ATTRIBUTE Unisphere-Cos-Shaping-Rate            ERX-VSA(177, string)

#JUNOS only
ATTRIBUTE Unisphere-Action-Reason               ERX-VSA(178, string)

#JUNOS only
ATTRIBUTE Unisphere-Service-Volume-Gigawords    ERX-VSA(179, integer) rt

#JUNOS only
ATTRIBUTE Unisphere-Update-Service              ERX-VSA(180, string) rt

#JUNOS only
# ERX-VSA(181) reserved for 15.1

```

```

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Remote-Id          ERX-VSA(182, string) c

#JUNOS only
ATTRIBUTE Unisphere-Acc-Loop-Encap              ERX-VSA(183, hexadecimal) c

#JUNOS only
ATTRIBUTE Unisphere-Inner-Vlan-Map-Id           ERX-VSA(184, integer) c

#JUNOS only
ATTRIBUTE Unisphere-Core-Facing-Interface       ERX-VSA(185, string) c

#JUNOS only
ATTRIBUTE Unisphere-Pcp-Port-Map               ERX-VSA(186, string) R

#JUNOS only
ATTRIBUTE Unisphere-vCPE-Lan-Extension          ERX-VSA(187, string) r

#JUNOS only
ATTRIBUTE Unisphere-vCPE-IPv4-Offload           ERX-VSA(188, string) r

#JUNOS only
ATTRIBUTE Jnpr-Input-Interface-Filter           ERX-VSA(191, string) r
ATTRIBUTE Jnpr-Output-Interface-Filter          ERX-VSA(192, string) r

#JUNOS only
ATTRIBUTE ERX-Bulk-CoA-Transaction-Id           ERX-VSA(194, integer) r
ATTRIBUTE ERX-Bulk-CoA-Identifler               ERX-VSA(195, integer) r

#JUNOS only
ATTRIBUTE Unisphere-IPv4-Input-Service-Set      ERX-VSA(196, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Set     ERX-VSA(197, string) r
ATTRIBUTE Unisphere-IPv4-Input-Service-Filter   ERX-VSA(198, string) r
ATTRIBUTE Unisphere-IPv4-Output-Service-Filter  ERX-VSA(199, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Set      ERX-VSA(200, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Set     ERX-VSA(201, string) r
ATTRIBUTE Unisphere-IPv6-Input-Service-Filter   ERX-VSA(202, string) r
ATTRIBUTE Unisphere-IPv6-Output-Service-Filter  ERX-VSA(203, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Profile-Name       ERX-VSA(204, string) r
ATTRIBUTE Unisphere-Adv-Pcef-Rule-Name          ERX-VSA(205, string) R

#JUNOS only
ATTRIBUTE ERX-DHCPv6-Options                    ERX-VSA(207, hexadecimal) CRO

```

```

ATTRIBUTE ERX-DHCP-Header          ERX-VSA(208, hexadecimal) c
ATTRIBUTE ERX-DHCPv6-Header        ERX-VSA(209, hexadecimal) c
ATTRIBUTE ERX-Acct-Request-Reason  ERX-VSA(210, integer) c
VALUE     ERX-Acct-Request-Reason  Acct-Start-Ack          1
VALUE     ERX-Acct-Request-Reason  Periodic-Interval-Interim 2
VALUE     ERX-Acct-Request-Reason  IP-Active              4
VALUE     ERX-Acct-Request-Reason  IP-Inactive              8
VALUE     ERX-Acct-Request-Reason  IPv6-Active             16
VALUE     ERX-Acct-Request-Reason  IPv6-Inactive            32
VALUE     ERX-Acct-Request-Reason  Session-Active           64
VALUE     ERX-Acct-Request-Reason  Session-Inactive         128
VALUE     ERX-Acct-Request-Reason  Line-Speed-Change        256
VALUE     ERX-Acct-Request-Reason  Address-Assignment-Change 512
VALUE     ERX-Acct-Request-Reason  CoA-Processing-Complete   1024

ATTRIBUTE ERX-Inner-Vlan-Tag-Protocol-Id  ERX-VSA(211, string) c

#JUNOS and JUNOSe
ATTRIBUTE Unisphere-Routing-Services      ERX-VSA(212, integer) r
VALUE     Unisphere-Routing-Services      disable 0
VALUE     Unisphere-Routing-Services      enable 1

#JUNOS only
ATTRIBUTE Unisphere-Interface-Set-Targeting-Weight ERX-VSA(213, integer) r
ATTRIBUTE Unisphere-Interface-Targeting-Weight    ERX-VSA(214, integer) r

#JUNOS and JUNOSe
##### tagged tunnel attributes

ATTRIBUTE Unisphere-Tunnel-VR-tag1          ERX-TAGGED-STRING-VSA(8, 1) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag1    ERX-TAGGED-STRING-VSA(9, 1) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag1    ERX-TAGGED-INT-VSA(33, 1) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag2          ERX-TAGGED-STRING-VSA(8, 2) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag2    ERX-TAGGED-STRING-VSA(9, 2) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag2    ERX-TAGGED-INT-VSA(33, 2) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag3          ERX-TAGGED-STRING-VSA(8, 3) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag3    ERX-TAGGED-STRING-VSA(9, 3) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag3    ERX-TAGGED-INT-VSA(33, 3) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag4          ERX-TAGGED-STRING-VSA(8, 4) rt

```


ATTRIBUTE Unisphere-Tunnel-Password-tag4	ERX-TAGGED-STRING-VSA(9, 4) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag4	ERX-TAGGED-INT-VSA(33, 4) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag5	ERX-TAGGED-STRING-VSA(8, 5) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag5	ERX-TAGGED-STRING-VSA(9, 5) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag5	ERX-TAGGED-INT-VSA(33, 5) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag6	ERX-TAGGED-STRING-VSA(8, 6) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag6	ERX-TAGGED-STRING-VSA(9, 6) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag6	ERX-TAGGED-INT-VSA(33, 6) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag7	ERX-TAGGED-STRING-VSA(8, 7) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag7	ERX-TAGGED-STRING-VSA(9, 7) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag7	ERX-TAGGED-INT-VSA(33, 7) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag8	ERX-TAGGED-STRING-VSA(8, 8) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag8	ERX-TAGGED-STRING-VSA(9, 8) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag8	ERX-TAGGED-INT-VSA(33, 8) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag9	ERX-TAGGED-STRING-VSA(8, 9) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag9	ERX-TAGGED-STRING-VSA(9, 9) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag9	ERX-TAGGED-INT-VSA(33, 9) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag10	ERX-TAGGED-STRING-VSA(8, 10) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag10	ERX-TAGGED-STRING-VSA(9, 10) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag10	ERX-TAGGED-INT-VSA(33, 10) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag11	ERX-TAGGED-STRING-VSA(8, 11) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag11	ERX-TAGGED-STRING-VSA(9, 11) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag11	ERX-TAGGED-INT-VSA(33, 11) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag12	ERX-TAGGED-STRING-VSA(8, 12) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag12	ERX-TAGGED-STRING-VSA(9, 12) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag12	ERX-TAGGED-INT-VSA(33, 12) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag13	ERX-TAGGED-STRING-VSA(8, 13) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag13	ERX-TAGGED-STRING-VSA(9, 13) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag13	ERX-TAGGED-INT-VSA(33, 13) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag14	ERX-TAGGED-STRING-VSA(8, 14) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag14	ERX-TAGGED-STRING-VSA(9, 14) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag14	ERX-TAGGED-INT-VSA(33, 14) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag15	ERX-TAGGED-STRING-VSA(8, 15) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag15	ERX-TAGGED-STRING-VSA(9, 15) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag15	ERX-TAGGED-INT-VSA(33, 15) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag16	ERX-TAGGED-STRING-VSA(8, 16) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag16	ERX-TAGGED-STRING-VSA(9, 16) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag16	ERX-TAGGED-INT-VSA(33, 16) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag17	ERX-TAGGED-STRING-VSA(8, 17) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag17	ERX-TAGGED-STRING-VSA(9, 17) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag17	ERX-TAGGED-INT-VSA(33, 17) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag18	ERX-TAGGED-STRING-VSA(8, 18) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag18	ERX-TAGGED-STRING-VSA(9, 18) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag18	ERX-TAGGED-INT-VSA(33, 18) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag19	ERX-TAGGED-STRING-VSA(8, 19) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag19	ERX-TAGGED-STRING-VSA(9, 19) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag19	ERX-TAGGED-INT-VSA(33, 19) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag20	ERX-TAGGED-STRING-VSA(8, 20) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag20	ERX-TAGGED-STRING-VSA(9, 20) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag20	ERX-TAGGED-INT-VSA(33, 20) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag21	ERX-TAGGED-STRING-VSA(8, 21) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag21	ERX-TAGGED-STRING-VSA(9, 21) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag21	ERX-TAGGED-INT-VSA(33, 21) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag22	ERX-TAGGED-STRING-VSA(8, 22) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag22	ERX-TAGGED-STRING-VSA(9, 22) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag22	ERX-TAGGED-INT-VSA(33, 22) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag23	ERX-TAGGED-STRING-VSA(8, 23) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag23	ERX-TAGGED-STRING-VSA(9, 23) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag23	ERX-TAGGED-INT-VSA(33, 23) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag24	ERX-TAGGED-STRING-VSA(8, 24) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag24	ERX-TAGGED-STRING-VSA(9, 24) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag24	ERX-TAGGED-INT-VSA(33, 24) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag25	ERX-TAGGED-STRING-VSA(8, 25) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag25	ERX-TAGGED-STRING-VSA(9, 25) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag25	ERX-TAGGED-INT-VSA(33, 25) rt

ATTRIBUTE Unisphere-Tunnel-VR-tag26	ERX-TAGGED-STRING-VSA(8, 26) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag26	ERX-TAGGED-STRING-VSA(9, 26) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag26	ERX-TAGGED-INT-VSA(33, 26) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag27	ERX-TAGGED-STRING-VSA(8, 27) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag27	ERX-TAGGED-STRING-VSA(9, 27) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag27	ERX-TAGGED-INT-VSA(33, 27) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag28	ERX-TAGGED-STRING-VSA(8, 28) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag28	ERX-TAGGED-STRING-VSA(9, 28) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag28	ERX-TAGGED-INT-VSA(33, 28) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag29	ERX-TAGGED-STRING-VSA(8, 29) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag29	ERX-TAGGED-STRING-VSA(9, 29) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag29	ERX-TAGGED-INT-VSA(33, 29) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag30	ERX-TAGGED-STRING-VSA(8, 30) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag30	ERX-TAGGED-STRING-VSA(9, 30) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag30	ERX-TAGGED-INT-VSA(33, 30) rt
ATTRIBUTE Unisphere-Tunnel-VR-tag31	ERX-TAGGED-STRING-VSA(8, 31) rt
ATTRIBUTE Unisphere-Tunnel-Password-tag31	ERX-TAGGED-STRING-VSA(9, 31) rt
ATTRIBUTE Unisphere-Tunnel-Sessions-tag31	ERX-TAGGED-INT-VSA(33, 31) rt

#JUNOS only

tagged service activation type attributes

ATTRIBUTE Unisphere-Service-Act-Type-tag1	ERX-TAGGED-INT-VSA(173, 1) rt
VALUE Unisphere-Service-Act-Type-tag1	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag1	op-script 2
ATTRIBUTE Unisphere-Service-Act-Type-tag2	ERX-TAGGED-INT-VSA(173, 2) rt
VALUE Unisphere-Service-Act-Type-tag2	dynamic-profile 1
VALUE Unisphere-Service-Act-Type-tag2	op-script 2

#####

unisphere.dct - Unisphere ERX dictionary

#####

2

PART

DHCP and DHCPv6 for Subscriber Management

- [DHCP for Subscriber Management | 1014](#)
-

DHCP for Subscriber Management

IN THIS CHAPTER

- [DHCP Overview | 1015](#)
- [DHCP Access Profiles for Subscriber Authentication and Accounting Parameters | 1029](#)
- [Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)
- [Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | 1047](#)
- [DHCP Options and Selective Traffic Processing | 1051](#)
- [Using DHCP Option 82 Information | 1078](#)
- [Default Services for DHCP Subscribers | 1091](#)
- [DHCP Client Attribute and Address Assignment | 1093](#)
- [DHCP Lease Times for IP Addresses | 1108](#)
- [DHCP Leasequery Methods | 1116](#)
- [DHCP Client Authentication With An External AAA Authentication Service | 1161](#)
- [Receiving DHCP Options From a RADIUS Server | 1166](#)
- [Common DHCP Configuration for Interface Groups and Server Groups | 1179](#)
- [Number of DHCP Clients Per Interface | 1189](#)
- [Maintaining DHCP Subscribers During Interface Delete Events | 1192](#)
- [Dynamic Reconfiguration of Clients From a DHCP Local Server | 1197](#)
- [Understanding Deferred NACK on DHCP Reconfigure Abort | 1206](#)
- [Conserving IP Addresses Using DHCP Auto Logout | 1208](#)
- [DHCP Short Cycle Protection | 1215](#)
- [DHCP Monitoring and Management | 1224](#)

DHCP Overview

IN THIS SECTION

- [Understanding Differences Between Legacy DHCP and Extended DHCP | 1015](#)
- [Extended DHCP Relay Agent Overview | 1019](#)
- [DHCP Relay Proxy Overview | 1021](#)
- [Minimum DHCP Relay Agent Configuration | 1023](#)
- [Example: DHCP Relay Agent Configuration with Multiple Clients and Servers | 1027](#)

Understanding Differences Between Legacy DHCP and Extended DHCP

IN THIS SECTION

- [New Features and Enhancements in Extended DHCP | 1015](#)
- [Benefits of Extended DHCP | 1017](#)
- [Change in Configuring DHCP Local Server in Extended DHCP Environment | 1017](#)
- [Legacy DHCP and Extended DHCP Server Hierarchy Levels Changes | 1018](#)

This topic covers the following sections:

New Features and Enhancements in Extended DHCP

Extended DHCP or JDHCP extends and enhances traditional DHCP operation. With the extended DHCP local server, the client configuration information resides in a centralized address-assignment pool, which supports advanced pool matching and address range selection. Any new features are only added to the Extended DHCP. Extended DHCP supports following features and enhancements:

- In extended DHCP, the address-assignment pools are external to the DHCP local server. The external address-assignment pools are managed by the **authd** process, independently of the DHCP local server, and can be shared by different client applications such as DHCP or PPPoE access. In legacy DHCP, client address pool and client configuration information reside on the DHCP server.

- Extended DHCP server interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide DHCP client authentication.
- You can configure the dynamic profile and authentication support on a global basis or for a specific group of interfaces.
- Extended DHCP local server supports IPv6 clients.
- Both DHCP local server and DHCPv6 local server support the specific address request feature, which enables you to assign a particular address to a client.
- The extended DHCP local server provides a minimal configuration to the DHCP client if the client does not have DHCP option 55 configured. The server provides the subnet mask of the address-assignment pool that is selected for the client. In addition to the subnet mask, the server provides the following values to the client if the information is configured in the selected address-assignment pool:
 - **router**—A router located on the client's subnet. This statement is the equivalent of DHCP option 3.
 - **domain name**—The name of the domain in which the client searches for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.
 - **domain name server**—A Domain Name System (DNS) name server that is available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.
- You can configure the local server to use DHCP option 82 information in the client PDU to determine which named address range to use for a particular client. The client configuration information, which is configured in the address-assignment pool, includes user-defined options, such as boot server, grace period, and lease time.
- The extended DHCP server supports following features:
 - *Graceful Routing Engine switchover* (GRES), which provides mirroring support for clients.
 - Virtual routing and forwarding (VRF). The extended DHCP is also referred to as virtual router (VR)-aware DHCP. See [Feature Explorer \(Virtual router aware DHCP\)](#) for a list of devices that support extended DHCP (VR-aware DHCP).

[Table 30 on page 1017](#) provides a comparison of the extended DHCP and a legacy DHCP configuration options.

Table 30: Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server

Feature	Legacy DHCP Local Server	Extended DHCP Local Server
Local address pools	X	X
External, centrally-managed address pools	–	X
Local configuration	X	X
External configuration using information from address-assignment pools or RADIUS servers	–	X
Dynamic-profile attachment	–	X
RADIUS-based subscriber authentication, and configuration using RADIUS attributes and Juniper Networks VSAs	–	X
IPv6 client support	–	X
Default minimum client configuration	X	X

Benefits of Extended DHCP

- Extended DHCP local server enhances traditional DHCP server operation by providing additional address assignment and client configuration functionality and flexibility in a subscriber-aware environment.
- Extended DHCP local server enables service providers to take advantage of external address-assignment pools and integrated RADIUS-based configuration capabilities in addition to the continued support of traditional local address pools.

Change in Configuring DHCP Local Server in Extended DHCP Environment

In extended DHCP, use the following steps to configure DHCP server and address assignment pool:

- Configure the extended DHCP local server on the device and specify how the DHCP local server determines which address-assignment pool to use.
- Configure the address-assignment pools used by the DHCP local server. The address-assignment pools contain the IP addresses, named address ranges, and configuration information for DHCP clients.

The extended DHCP local server and the address-assignment pools used by the server must be configured in the same logical system and routing instance.

Legacy DHCP and Extended DHCP Server Hierarchy Levels Changes

Legacy DHCP and extended DHCP servers can be configured at the hierarchy levels shown in [Table 31 on page 1018](#):

Table 31: Legacy DHCP and Extended DHCP Server Hierarchy Levels

DHCP Service	Hierarchy
Legacy DHCP server	<code>edit system services dhcp</code>
Extended DHCP server	<code>edit system services dhcp-local-server</code>
Legacy DHCP relay	<code>edit forwarding-options helpers bootp</code>
Extended DHCP relay	<code>edit forwarding-options dhcp-relay</code>
Legacy DHCP address pool	<code>edit system services dhcp pool</code>
Extended DHCP address pool	<code>edit access address-assignment pool</code>

Since legacy DHCP is deprecated, that is, the commands are 'hidden'. These commands do not show in the help nor automatic completion. When you use the option `show configuration` to display your configuration, the system displays the following warning:

```
##      ## Warning: configuration block ignored: unsupported platform (...)      ##
```

DHCP packets on non-configured interfaces are dropped

When you enable DHCP-Relay, the device inspects DHCP packets received on all interfaces. The interfaces that are not listed under the DHCP configuration are considered 'unconfigured'.

Depending on the configuration, DHCP packets received on unconfigured interfaces are dropped.

If the DHCP packets are dropped on 'unconfigured' interface, the DHCP traceoptions report it as:

```
May 25 18:26:31.796241 [MSTR][NOTE] [default:default][RLY][INET][irb.82] jdhcpd_packet_handle:
BOOTPREQUEST irb.82 arrived on unconfigured interface DISCOVER, flags 23, config 0x0
```

Some behaviors specific for some platforms have changed along the releases. See, [Release Notes](#).

Extended DHCP Relay Agent Overview

IN THIS SECTION

- [Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers | 1020](#)
- [DHCP Liveness Detection | 1021](#)

You can configure extended DHCP relay options on the router or on the switch and 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 DHCP client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

On the routers, you can use DHCP relay in carrier edge applications such as video/IPTV to obtain configuration parameters, including an IP address, for your subscribers.

On the switches, you can use DHCP relay to obtain configuration parameters including an IP address for DHCP clients.



NOTE: The extended DHCP relay agent options configured with the `dhcp-relay` statement are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, you cannot enable both the extended DHCP relay agent and the DHCP/BOOTP relay agent on the router at the same time.

For information about the DHCP/BOOTP relay agent, see [Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents](#).

You can also configure the extended DHCP relay agent to support IPv6 clients. See "[DHCPv6 Relay Agent Overview](#)" on page 1455 for information about the DHCPv6 relay agent feature.

To configure the extended DHCP relay agent on the router (or switch), include the `dhcp-relay` statement at the `[edit forwarding-options]` hierarchy level.

You can also include the `dhcp-relay` statement at the following hierarchy levels:

- `[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]`

Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers

The pattern of interaction among the DHCP Relay agent, DHCP client, and DHCP servers is the same regardless of whether the software installation is on a router or a switch. However, there are some differences in the details of usage.

On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, and the DHCP relay agent is configured on the router between the DHCP client and one or more DHCP servers.

On switches—In a typical network configuration, the DHCP client is on an access device such as a personal computer and the DHCP relay agent is configured on the switch between the DHCP client and one or more DHCP servers.

The following steps describe, at a high level, how the DHCP client, DHCP relay agent, and DHCP server interact in a configuration that includes two DHCP servers.

1. The DHCP client sends a discover packet to find a DHCP server in the network from which to obtain configuration parameters for the subscriber (or DHCP client), including an IP address.
2. The DHCP relay agent receives the discover packet and forwards copies to each of the two DHCP servers. The DHCP relay agent then creates an entry in its internal client table to keep track of the client's state.
3. In response to receiving the discover packet, each DHCP server sends an offer packet to the client. The DHCP relay agent receives the offer packets and forwards them to the DHCP client.

4. On receipt of the offer packets, the DHCP client selects the DHCP server from which to obtain configuration information. Typically, the client selects the server that offers the longest lease time on the IP address.
5. The DHCP client sends a request packet that specifies the DHCP server from which to obtain configuration information.
6. The DHCP relay agent receives the request packet and forwards copies to each of the two DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client's configuration parameters.
8. The DHCP relay agent receives the ACK packet and forwards it to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay agent installs a host route and Address Resolution Protocol (ARP) entry for this client.
11. After establishing the initial lease on the IP address, the DHCP client and the DHCP server use unicast transmission to negotiate lease renewal or release. The DHCP relay agent "snoops" on all of the packets unicast between the client and the server that pass through the router (or switch) to determine when the lease for this client has expired or been released. This process is referred to as *lease shadowing* or *passive snooping*.

DHCP Liveness Detection

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



NOTE: DHCP liveness detection either globally or per DHCP group.

DHCP Relay Proxy Overview

IN THIS SECTION

- [Benefits of Using DHCP Relay Proxy | 1022](#)
- [Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers | 1022](#)

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.

Normally, extended DHCP relay operates as a helper application for DHCP operations. 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.



NOTE: You cannot configure both DHCP relay proxy and extended DHCP local server on the same interface.

Benefits of Using DHCP Relay Proxy

DHCP relay proxy provides the following benefits:

- DHCP server isolation and DoS protection—DHCP clients are unable to detect the DHCP servers, learn DHCP server addresses, or determine the number of servers that are providing DHCP support. Server isolation also provides denial-of-service (DoS) protection for the DHCP servers.
- Multiple lease offer selection—DHCP relay proxy receives lease offers from multiple DHCP servers and selects a single offer to send to the DHCP client, thereby reducing traffic in the network. Currently, the DHCP relay proxy selects the first offer received.
- Support for both numbered and unnumbered Ethernet interfaces—For DHCP clients connected through Ethernet interfaces, when the DHCP client obtains an address, the DHCP relay proxy adds an access internal host route specifying that interface as the outbound interface. The route is automatically removed when the lease time expires or when the client releases the address. Note that DHCP Relay support for unnumbered Ethernet interfaces is not available on ACX7000 Devices (ACX7024, ACX7100, ACX7100, and ACX7509).
- Logical system support—DHCP relay proxy can be configured in a logical system, whereas a non-proxy mode DHCP relay cannot.

Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers

The DHCP relay agent is configured on the router (or switch), which operates between the DHCP client and one or more DHCP servers.

The following steps provide a high-level description of how DHCP relay proxy interacts with DHCP clients and DHCP servers.

1. The DHCP client sends a discover packet to locate a DHCP server in the network from which to obtain configuration parameters for the subscriber.
2. The DHCP relay proxy receives the discover packet from the DHCP client and forwards copies of the packet to each supporting DHCP server. The DHCP relay proxy then creates a client table entry to keep track of the client state.
3. In response to the discover packet, each DHCP server sends an offer packet to the client, which the DHCP relay proxy receives. The DHCP relay proxy does the following:
 - a. Selects the first offer received as the offer to sent to the client
 - b. Replaces the DHCP server address with the address of the DHCP relay proxy
 - c. Forwards the offer to the DHCP client.
4. The DHCP client receives the offer from the DHCP relay proxy.
5. The DHCP client sends a request packet that indicates the DHCP server from which to obtain configuration information—the request packet specifies the address of the DHCP relay proxy.
6. The DHCP relay proxy receives the request packet and forwards copies, which include the address of selected server, to all supporting DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client configuration parameters.
8. The DHCP relay proxy receives the ACK packet, replaces the DHCP server address with its own address, and forwards the packet to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay proxy installs a host route and Address Resolution Protocol (ARP) entry for the DHCP client.
11. After the initial DHCP lease is established, the DHCP relay proxy receives all lease renewals and lease releases from the DHCP client and forwards them to the DHCP server.

Minimum DHCP Relay Agent Configuration

IN THIS SECTION

- [Configuring IPv4 and IPv6 Addresses on the Loopback Interface | 1025](#)

This example shows the minimum configuration you need to use the extended DHCP relay agent on your Junos OS device. Ensure that the device can connect to the DHCP server.

In this example, you direct certain DHCP client traffic to a DHCP server. You specify an active server group to which each client groups traffic is forwarded. Add server IP addresses to the active server group. You can configure an interface group and specifying the DHCP relay interface for the group. The interface used as the DHCP relay agent can forward messages to specific servers.

Configure DHCP Option 82 and [forward-only](#) feature.

This example creates active server group named `my-dhcp-servers-group` with IP address `203.0.113.21`. The DHCP relay agent configuration is applied to a interfaces group named `my-dhcp-interfaces`. Within this group, the DHCP relay agent is enabled on interface `ge-0/0/1.0`.

1. Configure the option to forward the traffic, without creating a new subscriber session.

```
user@host# set forwarding-options dhcp-relay forward-only
```

2. Enable DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

```
user@host# set forwarding-options dhcp-relay relay-option-82 circuit-id use-interface-  
description device
```

Use the textual interface description instead of the interface identifier in the DHCP base option 82 Agent Circuit ID in DHCP packets that the DHCP relay agent sends to a DHCP server.

3. Configure DHCP server group and add the IP addresses of the DHCP server belonging to the group.

```
user@host# set forwarding-options dhcp-relay server-group my-dhcp-servers-group 203.0.113.2
```

4. Set the DHCP server group as active server group.

```
user@host# set forwarding-options dhcp-relay active-server-group my-dhcp-servers-group
```

The DHCP relay agent relays DHCP client requests to the DHCP servers defined in the active server group.

5. Configure an interface group and specify the DHCP relay interface for the group.

```
user@host# set forwarding-options dhcp-relay group my-dhcp-interf-group interface ge-0/0/1.0
```

DHCP relay runs on the interfaces defined in the group.



NOTE: To configure a switch with DHCP relay in forward-only mode, check whether your DHCP server supports DHCP Option 82. See [Verify support of Option-82 in DHCP Server](#) for details.

The forward-only option in DHCP relay configurations do not require the S-SA-FP license to be installed.

From configuration mode, confirm your configuration by entering the `show forwarding-options` command and verify your configuration.

```
user@srx-01# show forwarding-options
dhcp-relay {
  relay-option-82 {
    circuit-id {
      use-interface-description device;
    }
  }
  forward-only;
  server-group {
    my-dhcp-servers-group {
      203.0.113.21;
    }
  }
  active-server-group my-dhcp-servers-group;
  group my-dhcp-interf-group {
    interface ge-0/0/1.0;
  }
}
```

Configuring IPv4 and IPv6 Addresses on the Loopback Interface

When you have configured a DHCP server in a different service VRFs, you must configure IPv4 and IPv6 addresses on the loopback interface in the server VRF configuration for DHCP-relay function to work in all other VRFs.

Configure the dhcp-relay forward-only-replies option to enable DHCP response packets forwarded to the DHCP clients in the other VRF.

```
[edit routing-instances]
Svr-1 {
    instance-type vrf;
    routing-options {
        auto-export;
    }
    protocols {
        evpn {
            ip-prefix-routes {
                advertise direct-nexthop;
                encapsulation vxlan;
                vni 11000;
                export type5-export;
            }
        }
    }
    forwarding-options {
        dhcp-relay {
            dhcpv6 {
                forward-only-replies;
            }
            forward-only-replies;
        }
    }
    interface lo0.2;
    route-distinguisher 103.0.0.1:5000;
    vrf-import import-tenant;
    vrf-target target:5000:1;
    vrf-table-label;
}

lo0 {
    unit 0 {
        family inet {
            address 103.0.0.1/32;
        }
        family inet6 {
            address 1003::1/128;
        }
    }
}
```

```

    }
    unit 1 {
        family inet {
            address 103.0.0.1/32;
        }
        family inet6 {
            address 1003::1/128;
        }
    }
    unit 2 {
        family inet {
            address 103.0.0.2/32;
        }
        family inet6 {
            address 1003::2/128;
        }
    }
}

```

Example: DHCP Relay Agent Configuration with Multiple Clients and Servers

This example shows an extended DHCP relay agent configuration for a network that includes multiple DHCP clients and DHCP servers. Additional details follow the example.

```

[edit forwarding-options]
dhcp-relay {
    server-group {
        sp-1 {
            203.0.113.21;
            203.0.113.22;
        }
        sp-2 {
            203.0.113.31;
            203.0.113.32;
            203.0.113.33;
        }
    }
}
active-server-group sp-1;
overrides layer2-unicast-replies;
group clients_a {
    relay-option-82 circuit-id;
    interface fe-1/0/1.1;
    interface fe-1/0/1.2;
}

```

```

        interface fe-1/0/1.3;
    }
    group clients_b {
        relay-option-82 {
            circuit-id {
                prefix routing-instance-name;
            }
        }
        interface fe-1/0/1.4;
        interface fe-1/0/1.5;
        interface fe-1/0/1.6;
    }
    group eth_dslam_relay {
        active-server-group sp-2;
        overrides {
            trust-option-82;
            layer2-unicast-replies;
        }
        interface fe-1/0/1.7;
        interface fe-1/0/1.8;
        interface fe-1/0/1.9;
    }
}

```

This example creates two server-groups: sp-1, which includes DHCP server addresses 203.0.113.21 and 203.0.113.22, and sp-2, which includes DHCP server addresses 203.0.113.31, 203.0.113.32, and 203.0.113.33. The active server group to which the DHCP relay agent configuration applies is sp-1. A global override is set that causes the DHCP relay agent to use Layer 2 unicast transmission to send DHCP reply packets from the DHCP server to DHCP clients during the discovery process.

The example also creates three groups of subscribers and their associated Fast Ethernet interfaces: clients_a, clients_b, and eth_dslam_relay. These groups are configured to meet different needs, as follows:

- The clients_a and clients_b groups consist of basic subscribers. The service provider for these groups inserts option 82 information in the DHCP packets that are destined for the DHCP server.
- The subscribers in eth_dslam_relay are connected to an Ethernet digital subscriber line access multiplexer (DSLAM) that functions as a Layer 2 DHCP relay agent. The active server group for eth_dslam_relay is sp-2. Overrides are set for the eth_dslam_relay group that enable the DHCP relay agent to trust option 82 information and to use Layer 2 unicast transmission to send DHCP reply packets to DHCP clients during discovery.

RELATED DOCUMENTATION

[Address-Assignment Pools for Subscriber Management | 1476](#)

[DHCP Client Attribute and Address Assignment | 1093](#)

[DHCP Client Authentication With An External AAA Authentication Service | 1161](#)

[DHCP Monitoring and Management | 1224](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[DHCPv6 Monitoring and Management | 1464](#)

DHCP Access Profiles for Subscriber Authentication and Accounting Parameters

IN THIS SECTION

- [Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview | 1029](#)
- [Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | 1030](#)

Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview

Starting in Junos OS Release 14.2, access profiles enable you to specify subscriber access authentication and accounting parameters. After access profiles are created, you can attach them at the `[edit system services dhcp-local-server]` hierarchy level on a DHCP local server for DHCP or DHCPv6 subscribers and at the `[edit forwarding-options dhcp-relay]` hierarchy level on a DHCP relay agent for DHCP or DHCPv6 subscribers, group of subscribers, or group of interfaces.

If you configured a global access profile at the `[edit access profile profile-name]` hierarchy level for all DHCP or DHCPv6 clients on a router that functions as a DHCP local server or a DHCP relay agent, the access profile configured at the `[edit system services dhcp-local-server]` hierarchy level on a DHCP local server for DHCP or DHCPv6 subscribers and at the `[edit forwarding-options dhcp-relay]` hierarchy level on a DHCP relay agent for DHCP or DHCPv6 subscribers take precedence over the global access profile.

Configuring an access profile for DHCP subscribers at the DHCP relay agent level or the DHCP local server level provide you with the flexibility and effectiveness of enabling DHCP authentication and accounting for specific subscribers instead of enabling them at a global level. If no access profile is

configured at the DHCP relay agent level or the DHCP local server level, the global access profile becomes effective.

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

IN THIS SECTION

- [Attaching an Access Profile to All DHCP Subscriber or All DHCP Client Interfaces | 1030](#)
- [Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients | 1031](#)
- [Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces | 1032](#)

Starting in Junos OS Release 14.2, you can attach an access profile to a DHCP subscriber interface, to a DHCP client interface, to a group of subscriber interfaces, and to a specific subscriber or groups of subscribers. When a DHCP subscriber or DHCP client logs in, the specified access profile is instantiated and the services defined in the profile are applied to the interface, subscriber, or the group of interfaces or subscribers.

This topic contains the following sections:

Attaching an Access Profile to All DHCP Subscriber or All DHCP Client Interfaces

To attach an access profile to all DHCP subscribers or all DHCP client interfaces:

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set access-profile profile-name
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set access-profile profile-name
```

Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients

You use the group feature to group together a set of subscriber access profiles and then apply a common DHCP configuration to the named subscriber profile group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support groups.

Before you begin:

- Configure the group by entering the `group group-name` statement at the `[edit system services dhcp-local-server]` or the `[edit forwarding-options dhcp-relay]` hierarchy level. For DHCPv6 subscriber profiles, use the `dhcpv6` option at this hierarchy level.

To attach an access profile to a group of subscribers:

- At the DHCP configuration hierarchy, specify the name of the group and the access profile to attach to the group.
- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set group boston access-profile profile-name
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set group quebec access-profile profile-name
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set group boston access-profile profile-name
```

- For DHCPv6 relay agent:

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

Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces

Before you begin:

- Configure the interface group.

See ["Grouping Interfaces with Common DHCP Configurations" on page 1179](#).

To attach an access profile to a group of interfaces:

- At the DHCP configuration hierarchy, specify the name of the interface group and the access profile to attach to the group.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set group boston interface interface-name access-profile profile-name
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set group quebec interface interface-name access-profile profile-name
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set group boston interface interface-name access-profile profile-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set group quebec interface interface-name access-profile profile-name
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.2	Starting in Junos OS Release 14.2, access profiles enable you to specify subscriber access authentication and accounting parameters.
14.2	Starting in Junos OS Release 14.2, you can attach an access profile to a DHCP subscriber interface, to a DHCP client interface, to a group of subscriber interfaces, and to a specific subscriber or groups of subscribers.

RELATED DOCUMENTATION

[DHCP Overview](#) | **1015**

[DHCPv6 Local Server](#) | **1449**

[DHCPv6 Relay Agent](#) | **1455**

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings

IN THIS SECTION

- [Overriding the Default DHCP Local Server Configuration Settings](#) | **1034**
- [Overriding the Default DHCP Relay Configuration Settings](#) | **1036**
- [DHCP Behavior When Renegotiating While in Bound State](#) | **1039**
- [Sending Release Messages When Clients Are Deleted](#) | **1040**
- [Disabling Automatic Binding of Stray DHCP Requests](#) | **1041**
- [Enabling DHCP Relay Proxy Mode](#) | **1043**
- [Changing the Gateway IP Address \(giaddr\) Field to the giaddr of the DHCP Relay Agent](#) | **1044**
- [Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address](#) | **1044**
- [Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall](#) | **1045**
- [Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets](#) | **1045**

- [Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally | 1046](#)

Overriding the Default DHCP Local Server Configuration Settings

Subscriber management enables you to override certain default DHCP local server configuration settings. You can override the configuration settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP local server configuration options, include the `overrides` statement and its subordinate statements at the `[edit system services dhcp-local-server]` hierarchy level.
- To override DHCP local server configuration options for a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name]` hierarchy level.
- To override DHCP local server configuration options for a specific interface within a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name interface interface-name]` hierarchy level.
- To configure overrides for DHCPv6 local server at the global level, group level, or per-interface, use the corresponding statements at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

To override default DHCP local server configuration settings:

- (DHCPv4 and DHCPv6) Specify that you want to configure override options.
 - DHCPv4 overrides.

Global override:

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

Group-level override:

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

Per-interface override:

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

DHCPv6 overrides.

Global override:

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

Group level override:

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

Per-interface override:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit group group-name overrides interface interface-name
```

- (Optional) Override the maximum number of DHCP clients allowed per interface.
See ["Specifying the Maximum Number of DHCP Clients Per Interface" on page 1189](#).
- (Optional) Configure DHCP client auto logout.
See ["Automatically Logging Out DHCP Clients" on page 1210](#).
- (Optional) Enable processing of information requests from clients.
See ["Enabling Processing of Client Information Requests" on page 1105](#).
- (Optional) Specify that DHCP NAK and FORCERENEW messages support option 82 information.
See ["Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances" on page 1067](#).
- (Optional, DHCPv6 only) Specify a delegated pool name to use for DHCPv6 multiple address assignment.
See ["Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation" on page 1267](#).
- (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.

See ["Configuring DHCPv6 Rapid Commit \(MX Series, EX Series\)" on page 1450.](#)

- (Optional, DHCPv6 only) Specify that DHCPv6 local server return DNS server addresses as IA_NA or IA_PD suboptions rather than as a global DHCPv6 option.

See ["Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment" on page 1507.](#)

- (Optional, DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

See ["Automatically Logging Out DHCPv6 Clients" on page 1213.](#)

- (Optional) Specify that when the DHCP or DHCPv6 local server receives a Discover or Solicit message that has a client ID that matches the existing client entry, the local server deletes the existing client entry.

See ["DHCP Behavior When Renegotiating While in Bound State" on page 1039.](#)

- (Optional, DHCPv4 and DHCPv6) Specify that a short lease be sent to the client.

See ["Configuring DHCP Asymmetric Leasing" on page 1113.](#)

- (Optional, DHCPv4 and DHCPv6) Specify DHCP attributes globally or for groups.

See ["Configuring DHCP Attributes for All Clients or a Group of Clients" on page 1098.](#)

- Load balance traffic by allowing some local servers to respond to specific clients while preventing other local servers from responding immediately to these clients.

See ["Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers" on page 1047.](#)

Overriding the Default DHCP Relay Configuration Settings

You can override the default DHCP relay configuration settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP relay agent configuration options, include the overrides statement and its subordinate statements at the `[edit forwarding-options dhcp-relay]` hierarchy level.
- To override DHCP relay configuration options for a named group of interfaces, include the statements at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level.
- To override DHCP relay configuration options for a specific interface within a named group of interfaces, include the statements at the `[edit forwarding-options dhcp-relay group group-name interface interface-name]` hierarchy level.
- To configure overrides for DHCPv6 relay at the global level, group level, or per-interface, use the corresponding statements at the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level.

To override default DHCP relay agent configuration settings:

1. (DHCPv4 and DHCPv6) Specify that you want to configure override options.

- DHCPv4 overrides.

Global override:

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

Group-level override:

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

Per-interface override:

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

- DHCPv6 overrides.

Global override:

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

Group-level override:

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

Per-interface override:

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

2. (DHCPv4 only) Enable DHCP relay proxy mode.

See ["Enabling DHCP Relay Proxy Mode" on page 1043](#).

3. (DHCPv4 only) Overwrite the giaddr in DHCP packets that the DHCP relay agent forwards.

See ["Changing the Gateway IP Address \(giaddr\) Field to the giaddr of the DHCP Relay Agent" on page 1044.](#)

4. (DHCPv4 only) Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

See ["Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address" on page 1044.](#)

5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.

See ["Overriding Option 82 Information" on page 1078.](#)

6. (DHCPv4 only) Override the setting of the broadcast bit in DHCP request packets and use the Layer 2 unicast transmission method.

See ["Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets" on page 1045.](#)

7. (DHCPv4 only) Trust DHCP client packets that have a giaddr of 0 and that contain option 82 information.

See ["Enable Processing of Untrusted Packets So Option 82 Information Can Be Used" on page 1088.](#)

8. (DHCPv4 and DHCPv6) Override the maximum number of DHCP clients allowed per interface.

See ["Specifying the Maximum Number of DHCP Clients Per Interface" on page 1189.](#)

9. (DHCPv4 only) Configure client auto logout.

See ["DHCP Auto Logout Overview" on page 1208.](#)

10. (DHCPv4 and DHCPv6) Enable or disable support for DHCP snooped clients on interfaces.

See *Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent.*

11. (DHCPv4 and DHCPv6) Delay authentication of subscribers until the DHCP client sends a Request packet.

See the *delay-authentication*.

12. (DHCPv4 and DHCPv6) Send release messages to the DHCP server when clients are deleted.

See ["Sending Release Messages When Clients Are Deleted" on page 1040.](#)

13. (Optional) Specify that when the DHCP or DHCPv6 relay agent receives a Discover or Solicit message that has a client ID that matches the existing client entry, the relay agent deletes the existing client entry.

See ["DHCP Behavior When Renegotiating While in Bound State" on page 1039.](#)

14. (DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

See ["Automatically Logging Out DHCPv6 Clients" on page 1213.](#)

15. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.

See ["Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally" on page 1046.](#)

16. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.

See ["Disabling Automatic Binding of Stray DHCP Requests" on page 1041.](#)

17. (DHCPv4 and DHCPv6) Assign a single-session DHCP dual-stack group to a specified group of subscribers. You must assign the group to both legs of the DHCP dual stack.

See ["Configuring Single-Session DHCP Dual-Stack Support" on page 1320](#).

18. (Optional, DHCPv4 and DHCPv6) Specify that a short lease be sent to the client.

See ["Configuring DHCP Asymmetric Leasing" on page 1113](#).

DHCP Behavior When Renegotiating While in Bound State

All DHCP models (DHCPv4 and DHCPv6 local server and relay agent) use the same default behavior when receiving a DHCPv4 Discover or DHCPv6 Solicit message while in a bound state. In the default behavior, DHCP maintains the existing client entry when it receives a new Discover or Solicit message that has a client ID that matches the existing client. DHCP responds to the client with an Offer or Advertise message.

You can use the `delete-binding-on-renegotiation` statement to override the default behavior on DHCP local server or DHCP relay agent. You can configure the override on a global or group basis. In the override configuration, when DHCP is in a bound state and receives a Discover or Solicit message with a matching client entry, DHCP drops the message and does not process it. On a DHCP relay agent, the agent sends a Release message to the local server. DHCP cleans up the existing session and deletes the existing client entry, removing the binding. When a second Discover or Solicit message is received from the client, the message is processed and DHCP negotiation proceeds.



NOTE: In releases earlier than Junos OS Release 15.1, the default behavior for DHCPv6 local server and relay agent is the same as the override behavior in Junos OS Release 15.1 and later. For any release, the default behavior for DHCPv4 local server and relay agent is to maintain the existing client entry and respond without waiting for a second Discover or Solicit message.

For example, to configure DHCPv4 local server to override the default renegotiation behavior globally:

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

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

2. Specify that you want to configure an override action.

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

3. Specify that you want DHCP local server to override the default renegotiation behavior.

```
[edit system services dhcp-local-server overrides]
user@host# set delete-binding-on-renegotiation
```

For example, to configure DHCPv6 relay agent to override the default renegotiation behavior for an interface group:

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

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

2. Specify that the configuration is for an interface group.

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

3. Specify that you want to configure an override action.

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

4. Specify that you want DHCPv6 relay agent to override the default renegotiation behavior.

```
[edit forwarding-options dhcp-relay dhcpv6 group overrides]
user@host# set delete-binding-on-renegotiation
```

Sending Release Messages When Clients Are Deleted

By default, when DHCP relay and relay proxy delete a client, they do not send a release message to the DHCP server. You can override the default behavior and configure DHCP relay and relay proxy to send a release message whenever they delete a client. The release message sent by DHCP relay and relay proxy includes option 82 information.



NOTE: You must include the `send-release-on-delete` statement to configure DHCP relay and relay proxy to send the release message when the `client-discover-match` statement is included.

You can use the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level to override the default behavior for DHCPv6 relay agent.

To send a release message:

1. Specify that you want to configure override options.

- 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
```

2. Specify that you want DHCP relay and relay proxy (or DHCPv6 relay agent) to send a release message when clients are deleted.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set send-release-on-delete
```

Disabling Automatic Binding of Stray DHCP Requests

DHCP requests that are received but have no entry in the database are known as stray requests. By default, DHCP relay, DHCP relay proxy, and DHCPv6 relay agent attempt to bind the requesting client by creating a database entry and forwarding the request to the DHCP server. If the server responds with an ACK, the client is bound and the ACK is forwarded to the client. If the server responds with a NAK, the database entry is deleted and the NAK is forwarded to the client. This behavior occurs regardless of whether authentication is configured.

You can override the default configuration at the global level, for a named group of interfaces, or for a specific interface within a named group. Overriding the default causes DHCP relay, DHCP relay proxy, and DHCPv6 relay agent to drop all stray requests instead of attempting to bind the clients.



NOTE: Automatic binding of stray requests is enabled by default.

- To disable automatic binding behavior, include the `no-bind-on-request` statement when you configure DHCP overrides at the global, group, or interface level.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-bind-on-request
```

- To override the default behavior for DHCPv6 relay agent, configure the override at the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level.

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-bind-on-request
```

The following two examples show a configuration that disables automatic binding of stray requests for a group of interfaces and a configuration that disables automatic binding on a specific interface.

To disable automatic binding of stray requests on a group of interfaces:

1. Specify the named group.

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

2. Specify that you want to configure overrides.

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

3. Disable automatic binding for the group.

```
[edit forwarding-options dhcp-relay group boston overrides]
user@host# set no-bind-on-request
```

To disable automatic binding of stray requests on a specific interface:

1. Specify the named group of which the interface is a member.

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

2. Specify the interface on which you want to disable automatic binding.

```
[edit forwarding-options dhcp-relay group boston]
user@host# edit interface fe-1/0/1.2
```

3. Specify that you want to configure overrides.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2]
user@host# edit overrides
```

4. Disable automatic binding on the interface.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2 overrides]
user@host# set no-bind-on-request
```

Enabling DHCP Relay Proxy Mode

You can enable DHCP relay proxy mode on all interfaces or a group of interfaces.

To enable DHCP relay proxy mode:

1. Specify that you want to configure override options.

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

2. Enable DHCP relay proxy mode.

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

Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent

You can configure the DHCP relay agent to change the gateway IP address (giaddr) field in packets that it forwards between a DHCP client and a DHCP server.

To overwrite the giaddr of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server:

1. Specify that you want to configure override options.

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

2. Specify that the giaddr of DHCP packets is overwritten.

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

Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address

You can configure the DHCP relay agent to replace request and release packets with the gateway IP address (giaddr) before forwarding the packet to the DHCP server.

To replace the source address with giaddr:

1. Specify that you want to configure override options.

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

2. Specify that you want to replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

```
[edit forwarding-options dhcp-relay overrides]
user@host# set replace-ip-source-with giaddr
```

Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall

In network configurations where a firewall on the broadband network gateway (BNG) is between the DHCP relay agent and the DHCP server, only the BNG loopback address passes through the firewall. In that case, DHCP unicast packets are discarded. To enable DHCP unicast packets to pass through the BNG firewall, configure the source address in DHCP packets and DHCP messages to be the configured loopback address.

In addition to configuring the IP source address, on the DHCPv4 relay server, configure Link Selection (suboption 5) in option 82 information to cause the DHCP server to locate the correct address pool for the DHCP client when the server receives a forwarded packet, and Server ID Override (suboption 11) in option 82 information to set the server ID option in the DHCP packet.

To configure DHCPv4 relay agent to use the loopback address as the source address:

1. Configure the DHCPv4 relay agent to set the IP source address of DHCP packets to the configured loopback address.

```
[edit forwarding options dhcp-relay overrides]
user@host# set relay-source lo0
```

2. Configure the DHCPv4 relay agent to add Server ID and Link Selection suboptions to option 82 information:

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

To configure DHCPv6 relay agent to use the loopback address as the source address:

1. Configure the DHCPv6 relay agent to set the IP source address of DHCP packets to the configured loopback address.

```
[edit forwarding options dhcp-relay dhcpv6 overrides]
user@host# set relay-source lo0
```

Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets

You can configure the DHCP relay agent to override the setting of the broadcast bit in DHCP request packets. DHCP relay agent then instead uses the Layer 2 unicast transmission method to send DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.

To override the default setting of the broadcast bit in DHCP request packets:

1. Specify that you want to configure override options.

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

2. Specify that the DHCP relay agent uses the Layer 2 unicast transmission method.

```
[edit forwarding-options dhcp-relay overrides]  
user@host# set layer2-unicast-replies
```

Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally

You can disable DHCP relay on all interfaces or a group of interfaces.

To disable DHCP relay agent:

1. Specify that you want to configure override options.

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

2. Disable the DHCP relay agent.

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

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Common DHCP Configuration for Interface Groups and Server Groups | 1179](#)

[DHCP Options and Selective Traffic Processing | 1051](#)

[Using DHCP Option 82 Information | 1078](#)

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers

IN THIS SECTION

- [Load Balancing DHCP Local Servers by Delaying Responses to Clients | 1047](#)
- [Configuring a Delay in Local Server Response to DHCP Discover and Solicit Messages | 1048](#)

Load Balancing DHCP Local Servers by Delaying Responses to Clients

IN THIS SECTION

- [Benefits to Delaying DHCP Local Server Response | 1048](#)

In a network environment with multiple DHCP local servers and numerous DHCP clients, you might want to load balance traffic by allowing some local servers to respond to specific clients while preventing other local servers from responding immediately to these clients. Starting in Junos OS Release 16.1R1, you can configure a client-specific delay in response on DHCP local servers. When a DHCPv4 client sends a discover message or a DHCPv6 client sends a solicit message to the server network, all the corresponding (Same family) DHCP servers on the network receive the request at the same time, but servers that are configured with a delay do not respond to the client until the delay timer expires.

When the delay timer expires, the local server sends an offer or advertise message to the client. If the client is already bound, that means that a different server, one that has either no delay or a shorter delay, responded with an offer or advertise message to the client. In this case the server configured with the delay releases the client.

However, if the client does not receive a response from any server, it sends a second discover or solicit message. If the configured server receives the second message from the client before the original delay times out, it immediately sends a response to the client. This behavior enables the configured server to act as a redundant or back-up server for the server that was intended to handle the client.

[Table 32 on page 1048](#) lists the characteristics that you can use to identify DHCP clients for which responses are delayed and the corresponding DHCPv4 and DHCPv6 options you specify in the configuration.

Table 32: Characteristics to Identify Clients for Delayed Responses

Client Characteristic	DHCPv4	DHCPv6
Agent Circuit ID—A string that identifies the local circuit between the client and the DHCP relay agent, uniquely identifying the particular client.	Option 82, suboption 1	Option 18
Agent Remote ID—A string that uniquely identifies a client based on characteristics of the client, such as caller ID or user name.	Option 82, suboption 2	Option 37
User Class Identifier—A string representing a class or group to which the client belongs. For example, different user classes might identify a marketing group versus an accounting group.	Option 77	Option 15
Vendor Class Identifier—The IANA registered enterprise number for the vendor of the equipment running the client.	Option 60	Option 16

Benefits to Delaying DHCP Local Server Response

- Enable load to be distributed among many DHCP servers by causing certain clients to be preferably served by other servers.
- Enable redundancy among servers by allowing a server to respond in the event the preferred server does not.

Configuring a Delay in Local Server Response to DHCP Discover and Solicit Messages

You can configure a DHCPv4 or DHCPv6 local server to delay responding to discover and solicit messages, respectively, from clients. The server responds to the client only when the delay timer expires. You can configure the delay at global, group, and interface levels. To determine which clients are sent a delayed response, configure the server to identify specific hexadecimal or ASCII strings received in the message from the client. The local server compares the configured string with the value received DHCP options in the client message and delays the response depending on whether the received value matches the configured value, does not match it, or starts with the configured value.

To configure a delayed response to an offer message received from a DHCPv4 client:



NOTE: This procedure shows the global configuration. You can also configure the delay at the `[edit system services dhcp-local-server group group-name interface interface-name`

overrides] and [edit system services dhcp-local-server group *group-name* overrides] hierarchy levels.

1. Specify how long the DHCPv4 local server delays before responding to the client.

```
[edit system services dhcp-local-server overrides]
user@host# set delay-offer delay-time seconds
```

2. Specify the option received in the DHCPv4 offer message that identifies the client to receive a delayed response.

```
[edit system services dhcp-local-server overrides]
user@host# edit delay-offer based-on (option-60 | option-77 | option-82)
```

3. Specify how to match the received option.

- Match when the received ASCII or hexadecimal string is exactly the same as the configured string.

```
[edit system services dhcp-local-server overrides based-on (option-60 | option-77 |
option-82)]
user@host# set equals ascii ascii-string
user@host# set equals hexadecimal hexadecimal-string
```

- Match when the received ASCII or hexadecimal string is not exactly the same as configured string.

```
[edit system services dhcp-local-server overrides based-on (option-60 | option-77 |
option-82)]
user@host# set not-equals ascii ascii-string
user@host# set not-equals hexadecimal hexadecimal-string
```

- Match when the received ASCII or hexadecimal string starts with the configured string.

```
[edit system services dhcp-local-server overrides based-on (option-60 | option-77 |
option-82)]
user@host# set starts-with ascii ascii-string
user@host# set starts-with hexadecimal hexadecimal-string
```

To configure a delayed response to an advertise message received from a DHCPv6 client:



NOTE: This procedure shows the global configuration. You can also configure the delay at the [edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides] and [edit system services dhcp-local-server dhcpv6 group *group-name* overrides] hierarchy levels.

1. Specify how long the DHCPv6 local server delays before responding to the client.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set delay-advertise delay-time seconds
```

2. Specify the option received in the DHCPv6 advertise message that identifies the client to receive a delayed response.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# edit delay-advertise based-on (option-15 | option-16 | option-18 | option-37)
```

3. Specify how to match the received option.

- Match when the received ASCII or hexadecimal string is exactly the same as the configured string.

```
[edit system services dhcp-local-server dhcpv6 overrides based-on (option-15 | option-16 |
option-18 | option-37)]
user@host# set equals ascii ascii-string
user@host# set equals hexadecimal hexadecimal-string
```

- Match when the received ASCII or hexadecimal string is not exactly the same as configured string.

```
[edit system services dhcp-local-server dhcpv6 overrides based-on (option-15 | option-16 |
option-18 | option-37)]
user@host# set not-equals ascii ascii-string
user@host# set not-equals hexadecimal hexadecimal-string
```

- Match when the received ASCII or hexadecimal string starts with the configured string.

```
[edit system services dhcp-local-server dhcpv6 overrides based-on (option-15 | option-16 |
option-18 | option-37)]
```

```

user@host# set starts-with ascii ascii-string
user@host# set starts-with hexadecimal hexadecimal-string

```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
16.1R1	Starting in Junos OS Release 16.1R1, you can configure a client-specific delay in response on DHCP local servers.

RELATED DOCUMENTATION

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings 1033
Dynamic Reconfiguration of Clients From a DHCP Local Server 1197
DHCP Overview 1015
DHCPv6 Local Server 1449

DHCP Options and Selective Traffic Processing

IN THIS SECTION

- [DHCP Options and Selective Traffic Processing Overview | 1052](#)
- [Using DHCP Option Information to Selectively Process DHCP Client Traffic | 1054](#)
- [Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings | 1055](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings | 1055](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing | 1061](#)
- [DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | 1066](#)
- [Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | 1067](#)
- [DHCP-Initiated Service Change Based on Remote ID | 1071](#)
- [Configuring DHCP-Initiated Service Change Based on Remote ID | 1072](#)

DHCP Options and Selective Traffic Processing Overview

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

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

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

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

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

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

- DHCPv4 option 60 (Vendor Class Identifier)

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

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



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

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

Use the following match criteria to filter client traffic:

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



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

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

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



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

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

Using DHCP Option Information to Selectively Process DHCP Client Traffic

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings

To display the number of DHCP or DHCPv6 client packets that are dropped or forwarded during selective processing, use the following operational commands:

- `show dhcp relay statistics`
- `show dhcpv6 relay statistics`

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

IN THIS SECTION

- [Requirements | 1056](#)
- [Overview | 1056](#)

●	Configuration 1057
●	Verification 1059

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

Requirements

This example uses the following hardware and software components:

- MX Series 5G Universal Routing Platforms or EX Series Switches

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

- Configure DHCP relay agent.

See ["Extended DHCP Relay Agent Overview" on page 1019](#).

- (Optional) Configure a named DHCP local server group if you want to forward client traffic to a server group.

See ["Grouping Interfaces with Common DHCP Configurations" on page 1179](#).

Overview

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

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

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1057](#)
- [Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings | 1057](#)
- [Results | 1058](#)

To configure DHCP relay agent selective processing based on DHCP option information, 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 command into the CLI at the [edit] hierarchy level.

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

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

Step-by-Step Procedure

To configure DHCP relay selective processing:

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

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


2. Specify the DHCP option that DHCP relay agent uses to identify incoming client traffic.

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

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

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

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

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

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

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

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

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

Results

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

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

```

relay-option {
    option-number 60;
    equals {
        ascii video-gold {
            forward-only;
        }
    }
    equals {
        ascii video-bronze {
            local-server-group servergroup-15;
        }
    }
    default-action {
        drop;
    }
    starts-with {
        hexadecimal ffff {
            local-server-group servergroup-east;
        }
    }
}
}

```

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

Verification

IN THIS SECTION

- [Verifying the Status of DHCP Relay Agent Selective Traffic Processing | 1059](#)

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

Verifying the Status of DHCP Relay Agent Selective Traffic Processing

Purpose

Verify the DHCP relay agent selective traffic processing status.

Action

Display statistics for DHCP relay agent.

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

Packets dropped:

Total	30
Bad hardware address	1
Bad opcode	1
Bad options	3
Invalid server address	5
No available addresses	1
No interface match	2
No routing instance match	9
No valid local address	4
Packet too short	2
Read error	1
Send error	1
Option 60	1
Option 82	2

Messages received:

BOOTREQUEST	116
DHCPDECLINE	0
DHCPDISCOVER	11
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	105

Messages sent:

BOOTREPLY	0
DHCPOFFER	2
DHCPACK	1
DHCPNAK	0
DHCPFORCERENEW	0

Packets forwarded:

Total	4
BOOTREQUEST	2
BOOTREPLY	2

Meaning

The `Packets forwarded` field in the `show dhcp relay statistics` command output displays the number of client packets that have been forwarded as a result of the selective traffic processing configuration. In this example, the output indicates the total number of packets that DHCP relay agent has forwarded, as well as a breakdown for the number of `BOOTREQUEST` and `BOOTREPLY` packets forwarded.

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

IN THIS SECTION

- [Requirements | 1061](#)
- [Overview | 1062](#)
- [Configuration | 1062](#)
- [Verification | 1065](#)

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

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

Requirements

This example uses the following hardware and software components:

- MX Series 5G Universal Routing Platforms or PTX Series Packet Transport Routers

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

- Configure DHCPv6 relay agent.

See ["Extended DHCP Relay Agent Overview" on page 1019](#) and ["DHCPv6 Relay Agent Overview" on page 1455](#).

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

See ["Grouping Interfaces with Common DHCP Configurations" on page 1179](#).

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

See ["Grouping Interfaces with Common DHCP Configurations" on page 1179](#).

Overview

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

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

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1062](#)
- [Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings | 1063](#)
- [Results | 1064](#)

To configure group-level DHCPv6 relay agent selective processing based on DHCPv6 option information, 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 command into the CLI at the [edit] hierarchy level. The quick configuration assumes that the named interface group and the DHCP server groups have been previously configured.

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

Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings

Step-by-Step Procedure

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Results

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

```
dhcpv6 {
  group test-1 {
    relay-option {
      option-number 15;
      equals {
        ascii triple-gold {
          relay-server-group relayserver-triple-8;
        }
      }
    }
  }
}
```

```

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

```

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

Verification

IN THIS SECTION

- [Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing | 1065](#)

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

Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing

Purpose

Verify the DHCPv6 relay agent selective traffic processing status.

Action

Display statistics for DHCPv6 relay agent.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                0

Messages received:
    DHCPV6_DECLINE       0
    DHCPV6_SOLICIT       10
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE       0
    DHCPV6_REQUEST       10
    DHCPV6_CONFIRM       0
    DHCPV6_RENEW         0
    DHCPV6_REBIND        0
    DHCPV6_RELAY_REPL    0

Messages sent:
    DHCPV6_ADVERTISE     0
    DHCPV6_REPLY          0
    DHCPV6_RECONFIGURE    0
    DHCPV6_RELAY_FORW     0

Packets forwarded:
    Total                4
    FWD REQUEST          2
    FWD REPLY            2
```

Meaning

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

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs

In some service provider networks, the service network in which the DHCP server resides is isolated from the actual subscriber network. This separation of the service and subscriber networks can sometimes introduce potential security issues, such as route leaking.

Starting in Junos OS Release 14.2, you can use the DHCP relay agent to provide additional security when exchanging DHCP messages between different virtual routing and forwarding instances (VRFs). The DHCP relay agent can ensure that there is no direct routing between the client VRF and the DHCP server VRF, and that only acceptable DHCP packets are relayed across the two VRFs. Subscriber management supports the cross-VRF message exchange for both DHCP and DHCPv6 packets.

To exchange DHCP messages between different VRFs, you must enable both the server-side and the client-side of the DHCP relay agent to recognize and forward acceptable traffic based on DHCP option information in the packets. The message exchange uses the following DHCP options to identify the traffic to be relayed.

- Agent Circuit ID (DHCP option 82 suboption 1) for DHCPv4 packets
- Relay Agent Interface-ID (DHCPv6 option 18) for DHCPv6 packets

Statistics for DHCP packets using the cross-VRF message exchange are counted in the client VRF.

The following list describe how DHCP relay agent exchanges messages between the DHCP clients and DHCP server in different VRFs:

- Packets from DHCP client to DHCP server—DHCP relay agent receives the DHCP packet from the client in the client VRF, and then inserts the appropriate DHCP option 82 suboption 1 or DHCPv6 option 18 attribute into the packet. The relay agent then forwards the packet to the DHCP server in the server's VRF.
- Packets from DHCP server to DHCP client—DHCP relay agent receives the DHCP reply message from the DHCP server in the server VRF. The relay agent derives the client's interface, including VRF, from the DHCP option 82 suboption 1 or DHCPv6 option 18 attribute in the packet in the DHCP server VRF. The relay agent then forwards the reply message to the DHCP client in the client's VRF.

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances

IN THIS SECTION

- [Client-Side Support | 1069](#)
- [Server-Side Support | 1069](#)
- [DHCP Local Server Support | 1070](#)

Starting in Junos OS Release 14.2, you can configure DHCP relay agent to provide additional security when exchanging DHCP messages between a DHCP server and DHCP clients that reside in different virtual routing and forwarding instances (VRFs).

You can configure DHCP relay agent to provide additional security when exchanging DHCP messages between a DHCP server and DHCP clients that reside in different virtual routing instances. This type of configuration is for a *stateless* DHCP relay connection between a DHCP server and a DHCP client, when the DHCP server resides in a network that must be isolated from the client network.

A stateless DHCP relay agent does not maintain dynamic state information about the DHCP clients and does not maintain a static route for the traffic to flow between the client and server routing instances.

To enable the DHCP message exchange between the two VRFs, you configure each side of the DHCP relay to recognize and forward acceptable traffic based on the DHCP option information in the packets. The acceptable traffic is identified by either the Agent Circuit ID (DHCP option 82 suboption 1) for DHCPv4 packets or the Relay Agent Interface-ID (DHCPv6 option 18) for DHCPv6 packets.

The following list provides an overview of the tasks required to create the DHCP message exchange between the different VRFs:

- Client-side support—Configure the DHCP relay agent `forward-only` statement to specify the VRF location of the DHCP server, to which the DHCP relay agent forwards the client packets with the appropriate DHCP option information. The `forward-only` statement ensures that DHCP relay agent does not create a new session or perform any other subscriber management operations (such as creating dynamic interfaces or maintaining leases).

You can optionally configure a specific logical system and routing instance for the server VRF. If you do not specify a logical system or routing instance, then DHCP uses the local logical system and routing instance from which the configuration is added.

- Server-side support—Configure the DHCP relay agent `forward-only-replies` statement so the DHCP relay agent forwards the reply packets that have the appropriate DHCP option information. This statement also ensures that DHCP relay agent does not create a new session or perform any other subscriber management operations.



NOTE: You do not need to configure the `forward-only-replies` statement if the DHCP client and DHCP server reside in the same logical system/routing instance.

- DHCP local server support—Configure the DHCP local server to support option 82 information in DHCP NAK and `forcerenew` messages. By default, the two message types do not support option 82.
- Additional support—Ensure that the following required support is configured:
 - Proxy ARP support must be enabled on the server-facing interface in the DHCP server VRF so that the DHCP relay agent can receive and respond to the ARP requests for clients and the client-facing interface in the DHCP server VRF.
 - Routes must be available to receive the DHCP packets from the DHCP server in the server VRF for the clients reachable in the client VRF.

The following procedures describe the configuration tasks for creating the DHCP message exchange between the DHCP server and clients in different VRFs.

Client-Side Support

To configure support on the client side of the DHCP relay agent:

1. Enable DHCP relay agent configuration.

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

2. Specify the DHCP server VRF to which the DHCP relay agent forwards the packets from the DHCP client. DHCP relay agent forwards the acceptable packets that have the appropriate DHCP option information, but does not perform any additional subscriber management operations. You can configure the `forward-only` statement globally or for a named group of interfaces, and for DHCPv4 or DHCPv6. You can specify the current, default, or a specific logical system or routing instance for the server VRF.

The following example configures the `forward-only` statement globally for DHCPv4, and specifies the default logical system and routing instance:

```
[edit forwarding-options dhcp-relay]
user@host# set forward-only logical-system default routing-instance default
```



NOTE: For local DHCPv4 clients, the DHCP relay agent adds the Agent Circuit ID option. However, if the Agent Circuit ID option is already present in the packet, you must ensure that the DHCP server supports the option 82 Vendor-Specific Information suboption (suboption 9).

If the `forward-only` statement is configured at the `[edit forwarding-options dhcp-relay relay-option]` hierarchy level, then that relay-option action takes precedence over the configuration of the `forward-only` statement for the DHCP cross-VRF message exchange.

Server-Side Support

To configure the cross-VRF message exchange support on the server side of the DHCP relay:



NOTE: You do not need to configure the `forward-only-replies` statement if the DHCP client and DHCP server reside in the same logical system/routing instance.

1. Enable DHCP relay agent configuration.

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

2. Configure the DHCP relay agent to forward the DHCP packets from the DHCP server VRF to the client. DHCP relay agent only forwards the packets, and does not perform any additional subscriber management operations. You can configure the forward-only-replies statement globally for DHCPv4 and DHCPv6.

The following example configures the forward-only-replies statement globally for DHCPv4.

```
[edit forwarding-options dhcp-relay]
user@host# set forward-only-replies
```

DHCP Local Server Support

To configure the DHCP local server to support option 82 information in NAK and forcerenew messages; the cross-VRF message exchange feature uses the option 82 or DHCPv6 option 18 information to determine the client VRF:

1. Enable DHCP local server configuration.

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

2. Specify that you want to configure an override option.

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

3. Configure DHCP local server to override the default behavior and support option 82 information in DHCP NAK and forcerenew messages. You can configure the override action globally, for a group of interfaces, or for a specific interface.

```
[edit system services dhcp-local-server overrides]
user@host# set include-option-82 forcerenew nak
```

DHCP-Initiated Service Change Based on Remote ID

IN THIS SECTION

- [Benefits of DHCP-Initiated Service Change Based on Remote ID | 1072](#)

Subscriber management enables you to update a DHCP client's current service through the use of the client's remote ID (Agent Remote ID). The remote ID can be in option 82, suboption 2 for DHCPv4 clients, or option 37 for DHCPv6 clients.

When a DHCP client is initially established, DHCP preserves the client's incoming remote ID in the DHCP client database. When receiving a rebind or renew message for that client, DHCP compares the client's initial remote ID to the remote ID in the DHCP renew or rebind message. If the two remote IDs do not match, DHCP local server tears down the existing binding and sends a NAK message (or logical NAK for DHCPv6), which causes the client to initiate a reconnect sequence. When the client reconnects, the DHCP local server activates the new service, which is encoded within the new Agent Remote ID string.

You can configure the router to support the remote ID service change feature globally or for a specific group, and you can configure the support on DHCP local server and DHCP relay agent.

In a dual-stack environment, the DHCP-initiated service change feature requires that a client's DHCPv4 and DHCPv6 sessions reside over the same VLAN (1:1 mapping) and that the Agent Remote ID strings for the two sessions are identical. The dual-stack support also requires that the same dynamic client profile is applied to both the DHCPv4 and DHCPv6 networks, to ensure remote ID consistency between the two networks. When DHCP detects a remote ID mismatch in one session of the dual stack, DHCP tears down that session. The incoming remote ID is then compared to the other session of the dual stack, and if there is a mismatch, that other session is torn down gracefully.

During the graceful teardown process, if the other session is currently in the bound state, that session then transitions to the deferred delete state. The deferred delete state allows the session that detected the change to be reestablished immediately with the new service plan, while enabling the router to gracefully tear down the other session by sending NAK messages in response to the subsequent renew and rebind messages.

As part of the DHCP-initiated service change feature, AAA can set a session's client profile. AAA obtains the client profile from the remote-ID, and writes the profile into the session database. A client profile that AAA writes into the database always takes precedence over any local DHCP configuration.

A change in the Agent Remote ID can also initiate a service change during reauthentication. You cannot configure both the `remote-id-mismatch` statement and the `reauthenticate` statement at the global level, [`edit system services dhcp-local-server`]. However, DHCP precedence rules do permit you to configure both

statements when they are at different levels. For example, you can configure `reauthenticate` at the global level and `remote-id-mismatch` for DHCPv6 at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level or for a specific group at the `[edit system services dhcp-local-server group name]` hierarchy level, and so on.

Benefits of DHCP-Initiated Service Change Based on Remote ID

- Enables DHCP server or relay agent to update the client's service when based on the remote ID when the client remote ID changes and does not match the previously stored agent remote ID. In a dual-stack environment, enables the DHCP server or relay agent to help ensure the remote ID consistency between sessions of the dual stack.

Configuring DHCP-Initiated Service Change Based on Remote ID

This topic describes how to configure support for DHCP-initiated service change on the DHCP local server and the DHCP relay agent.

Configuring DHCP local server

You can configure support for DHCP-initiated service change for DHCP local server and DHCPv6 local server globally, or for a named group of interfaces.

To configure DHCP local server to support DHCP-initiated service change for a named group:

1. Before starting the configuration for the DHCP-initiated service change feature, ensure that the DHCP or DHCPv6 relay agent is configured to override the default behavior and send a release message to the DHCP server when a remote ID mismatch occurs. This configuration is required because the relay agent cannot directly tear down client bindings; the release packet signals the DHCP server to tear down the original binding.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set overrides send-release-on-delete
```

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

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

3. Specify the named group that you want to configure.

```
[edit system services dhcp-local-server]
user@host# edit group northwest-321
```

4. Specify that, for the named group, DHCP local server matches the initial and new client remote IDs, and then performs the `disconnect` action when a mismatch occurs.

```
[edit system services dhcp-local-server group northwest-321]  
user@host# set remote-id-mismatch disconnect
```

When a mismatch occurs, DHCP local server tears down the existing binding and sends a NAK message to the client, which initiates the client reconnect sequence. The new service, which is encoded in the Agent Remote ID string, is then activated when the client reconnects.

Configuring DHCP relay agent

You can configure support for DHCP-initiated service change for DHCP relay agent and DHCPv6 relay agent globally, or for a named group of interfaces.

The following example shows the steps to configure DHCPv6 relay agent to support DHCP-initiated service change on a global basis.

1. Before starting the configuration for the DHCP-initiated service change feature, ensure that the DHCPv6 relay agent is configured to override the default behavior and send a release message to the DHCP server when a remote ID mismatch occurs. This configuration is required because the relay agent cannot directly tear down client bindings; the release packet signals the DHCP server to tear down the original binding.

```
[edit forwarding-options dhcp-relay dhcpv6]  
user@host# set dhcp-relay overrides send-release-on-delete
```

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

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

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

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


4. Specify that DHCPv6 relay agent matches the initial and new client remote IDs, and then performs the disconnect action when a mismatch occurs.

```
[edit system services dhcp-local-server group northwest-321]
user@host# set remote-id-mismatch disconnect
```

When a mismatch occurs, DHCPv6 relay agent sends the release message to the DHCPv6 local server and a logical NAK message (a reply packet with a 0 lifetime) to the client. The server then tears down the existing binding, and the client initiates the reconnect sequence. The new service, which is encoded in the Agent Remote ID string, is activated when the client reconnects.

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

IN THIS SECTION

- [Processing of DHCPv4 and DHCPv6 Destination Addresses | 1075](#)
- [Processing Order and Actions | 1076](#)
- [Benefits of DHCP Relay Forward-Only Action | 1077](#)

The DHCP relay agent entry, which can be created on a DHCPv4 or DHCPv6 server, is useful for authentication, authorization, accounting, applying filtering, ensuring quality of service (QoS) on the client, and processing of options specified in the packet. The creation of the relay agent or client entry involves participation of the `jdhcpd` process memory resources, session database resources, authentication procedure, accounting, dynamic profile instantiation, dynamic interface creation, firewall, class-of-service (CoS)-association, and more. Customer networks can contain non-customer controlled bindings for which they might not want these relay agent entry functionalities. When a customer's network has such traffic, creation of relay agent entries—which is related to client entry creation—unnecessarily utilizes resources and can also result in wrong association of profiles, because in the current network scenario, all the traffic received from a specific interface is forwarded, without processing any destination address.

Starting in Junos OS Release 17.4R1, a forward-only configuration can be enabled on the broadband network gateway (BNG) device for non-customer traffic along with unknown DHCP server address. The configuration of the `forward-only` statement, along with the new DHCP options—`option-54` for DHCPv4 and `option-2` for DHCPv6, avoids the creation of the DHCP relay agent entry on the BNG and ensures that traffic is forwarded to the specified destination address.

With these configurations, administrators are able to determine to which servers the clients are bound; which of the clients need to have a relay client entry created and dynamic profile and policies applied, and so on; and for whom (non-customer) the forward-only configuration is enabled.

The two new configuration statements—options-54 and options-2—are introduced for processing destination addresses.

Processing of DHCPv4 and DHCPv6 Destination Addresses

Administrators can configure the forward-only statement to avoid the creation of non-customer client entries. The `jdhcpd` process compares the server identifier (option-54 for DHCPv4 and option-2 for DHCPv6) with or without destination address in the incoming packet along with the configured server address. If the server identifier and configured server address match, the action is to only forward without creating the client entry.

On nonpassive relay, configuration of the `server-match` statement means implicit enabling of the `delay-authentication` statement for the clients for which the `server-match` statement is processed. You can also configure options 60 and 77 (for DHCPv4) or options 15 and 16 (for DHCPv6) optionally together with associated processing. Configuration of these options also includes the specification of the order in which they are processed. If these options are not configured, the default order is 60, 77 for DHCPv4 and 15, 16 for DHCPv6.



NOTE: DHCPv6 option-16 data as defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, comprises a 4-byte enterprise number and the variable length vendor-class data. The enterprise number is a number registered by the vendor with IANA. As such, it is anticipated that configuring an ASCII match in conjunction with option-16 relay-option match might not work because the enterprise number must coincide with the value of a printable ASCII character.

A similar restriction exists for DHCPv4 option-77 statement because the option-77 data might be subdivided to include suboptions and sublengths. Because of this, configuring an ASCII match with option-77 relay-option match might not work.

On nonpassive relay, if a request packet is received in the *rebind* phase and the corresponding relay entry is not present, then you need to first configure the `bind-on-request` statement, following which the relay entry is created and the packet is forwarded. After an acknowledgment is received, the `jdhcpd` process verifies the source address (with or without option 54 for DHCPv4 and option 2 for DHCPv6) within the `server-match` configuration. If the verification results in a stateless entry, then the relay entry is deleted.

The administrator can specify the DHCP unique identifier (DUID) with or without the address of a server. The `jdhcpd` process first processes address statements and then the DUID statements. If the same server address is specified in address statements and the DUID statement, then it is the administrator's responsibility to specify the same action for both address and the DUID statements.

On both passive and nonpassive relays, if the received packet contains a relay forward header and the destination address is multicast or a link-local address, then the packet is forwarded without any further processing.



NOTE: For both DHCPv4 and DHCPv6 subscribers, the `relay-option` and `server-match` statements are at the same hierarchy and have the same priority.

Processing Order and Actions

Relay options and server match processing are mutually exclusive. Although they are at the same hierarchy and have the same priority, for implementation purpose, you process relay options followed by server match.

Following are the DHCPv4 relay option processing actions:

- `drop`—Discard when there is a match.
- `forward-only`—Forward without client services, when there is a match.
- `local-server-group`—Name of DHCP local server group when there is a match.
- `relay-server-group`—Name of DHCP relay server group when there is a match.

Following are the DHCPv6 relay option processing actions:

- `drop`—Discard when there is a match.
- `forward-only`—Forward without client services, when there is a match.
- `relay-server-group`—Name of DHCP relay server group when there is a match.



NOTE: The DHCPv6 server match for IPv6 address is available in passive relay only.

The default and configurable option orders are processed as shown in [Figure 4 on page 1077](#). If you need the option order to be reversed (either DHCPv4 or DHCPv6), then configure `option-order 77,66` statement for DHCPv4 or `option-order 16,15` statement for DHCPv6 at the `[edit forwarding-options dhcp-relay relay-option]` hierarchy level.

Figure 4: DHCPv4 and DHCPv6 Option Order

[DHCPv4](#)
Default order:
option-60 (equals -> starts-with) -> option-77 (equals -> start-with) -> option-60 (default-action) -> option-77 (default-action) -> not-present

When option-order is configured as 77,60 in CLI:
option-77 (equals -> starts-with) -> option-60 (equals -> starts-with) -> option-77 (default-action) -> option-60 (default-action) -> not-present

[DHCPv6](#)
Default order:
option-15 (equals -> starts-with) -> option-16 (equals -> start-with) -> option-15 (default-action) -> option-16 (default-action) -> not-present

When option-order is configured as 16,15 in CLI:
option-16 (equals -> starts-with) -> option-15 (equals -> starts-with) -> option-16 (default-action) -> option-15 (default-action) -> not-present

g2002000

Benefits of DHCP Relay Forward-Only Action

- Reduces the unnecessary consumption of resources for non-customer controlled bindings that do not need the relay agent entries made when client entries are created.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.1	Starting in Junos OS Release 15.1, the selective traffic processing feature lets you manage multivendor networks with the extended DHCP and DHCPv6 relay agent.
15.1	Starting in Junos OS Release 15.1, you can configure the DHCP relay agent to selectively process client traffic.
14.2	Starting in Junos OS Release 14.2, you can use the DHCP relay agent to provide additional security when exchanging DHCP messages between different virtual routing and forwarding instances (VRFs).
14.2	Starting in Junos OS Release 14.2, you can configure DHCP relay agent to provide additional security when exchanging DHCP messages between a DHCP server and DHCP clients that reside in different virtual routing and forwarding instances (VRFs).

RELATED DOCUMENTATION

| [DHCP Overview](#) | 1015

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[Common DHCP Configuration for Interface Groups and Server Groups | 1179](#)

[RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 727](#)

Using DHCP Option 82 Information

IN THIS SECTION

- [Using DHCP Relay Agent Option 82 Information | 1078](#)
- [Enable Processing of Untrusted Packets So Option 82 Information Can Be Used | 1088](#)
- [Extracting an Option 82 or Option 37 Substring to Create an Interface Set | 1089](#)

Using DHCP Relay Agent Option 82 Information

IN THIS SECTION

- [Configuring Option 82 Information | 1079](#)
- [Overriding Option 82 Information | 1082](#)
- [Including a Prefix in DHCP Options | 1083](#)
- [Including a Textual Description in DHCP Options | 1086](#)

Subscriber management enables you to configure the DHCP relay agent to include additional option 82 information in the DHCP packets that the relay agent receives from clients and forwards to a DHCP server. The DHCP server uses the additional information to determine the IP address to assign to the client. The server might also use the information for other purposes—for example, to determine which services to grant the client, or to provide additional security against threats such as address spoofing. The DHCP server sends its reply back to the DHCP relay agent, and the agent removes the option 82 information from the message and forwards the packet to the client.

To configure support for the DHCP relay agent information option 82, you use the `relay-option-82` statement. You can configure the DHCP relay agent to include the following suboptions in the packet the relay agent sends to the DHCP server:

- Agent Circuit ID (suboption 1)—An ASCII string that identifies the interface on which the client DHCP packet is received.



NOTE: If `relay-option-82` is configured, but none of the attributes under `relay-option-82` (that is, `circuit-id` | `remote-id` | `server-id-override`) are explicitly configured, then the default behavior is for the `circuit-id` (that is, suboption 1) to always be included in the option-82 value. This is true whether or not the vendor-specific attribute under `relay-option-82` is configured.

- Agent Remote ID (suboption 2)—An ASCII string assigned by the DHCP relay agent that securely identifies the client.

You can configure the option 82 support globally or for a named group of interfaces.

To restore the default behavior, in which option 82 information is not inserted into DHCP packets, you use the `delete relay-option-82` statement.



NOTE: The DHCPv6 relay agent provides similar Agent Circuit ID and Agent Remote ID support for DHCPv6 clients. For DHCPv6, subscriber management uses DHCPv6 option 18 to include the circuit ID in the packets that the relay agent sends to a DHCPv6 server, and option 37 to include the remote ID in the packets. See ["DHCPv6 Relay Agent Options" on page 1456](#).

The following sections describe the option 82 operations you can configure:

Configuring Option 82 Information

You use the `relay-option-82` statement to configure the DHCP relay agent to insert option 82 information in DHCP packets that the relay agent receives from clients and forwards to a DHCP server. When you configure option 82, you can include one of the suboption statements to specify the type of information you want to include in the DHCP packets. If you configure option 82 without including one of the suboption statements, the Agent Circuit ID option is included by default. Use the `circuit-id` statement to include the Agent Circuit ID (suboption 1) in the packets, or the `remote-id` statement to include the Agent Remote ID (suboption 2).

You can optionally configure DHCP relay agent to include a prefix or the interface description as part of the suboption information. If you specify the `circuit-id` or `remote-id` statement without including any of the optional `prefix`, `use-interface-description`, `use-vlan-id`, `include-irb-and-l2`, or `no-vlan-interface-name`

statements, the format of the Agent Circuit ID or Agent Remote ID information for Fast Ethernet (fe), Gigabit Ethernet (ge), and integrated routing and bridging (irb) interfaces is one of the following, depending on your network configuration:

- For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs, stacked VLANs (S-VLANs), or bridge domains:

```
(fe | ge)-fpc/pic/port.subunit
```



NOTE: For remote systems, the *subunit* is required and is used to differentiate an interface.

- For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs:

```
(fe | ge)-fpc/pic/port:vlan-id
```

- For Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs:

```
(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```



NOTE: Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridging domain that has a Layer 3 protocol configured.

The interface to bridge domain relationship might be implicit (the interface is mapped to the bridge domain by the system based on the VLAN tag) or explicit (the interface is mapped to the bridge domain by configuring it in the bridge domain definition). For the explicit case, tagging might not be relevant for the mapping.

In the case of an IRB interface, the format displays the Layer 2 interface instead of the IRB interface along with the bridge domain name. For IRB interfaces (or other pseudo devices) the default format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

```
(fe | ge)-fpc/pic/port.subunit:bridge-domain-name
```

- IRB interfaces that use VLANs:

```
(fe | ge)-fpc/pic/port.subunit:vlan-name
```

To include the IRB interface name with the Layer 2 interface name, configure the `include-irb-and-l2` statement. The format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

```
(fe | ge)-fpc/pic/port:bridge-domain-name+irb.subunit
```

- IRB interfaces that use VLANs:

```
(fe | ge)-fpc/pic/port:vlan-name+irb.subunit
```

To include only the IRB interface name without the Layer 2 interface and bridge domain or VLAN, configure the `no-vlan-interface-name` statement. The format is as follows:

```
irb.subunit
```

To enable insertion of option 82 information:

1. Specify that you want to configure option 82 support.

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

2. Configure the DHCP relay agent to insert the Agent Circuit ID suboption, the Agent Remote ID suboption, or both.

- To insert the Agent Circuit ID:

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


- To insert the Agent Remote ID:

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

- To insert both, configure both set commands.
3. (Optional) Configure a prefix that is used in the option 82 information in the DHCP packets.
See ["Including a Prefix in DHCP Options" on page 1078](#).
 4. (Optional) Configure the DHCP relay agent to include the interface's textual description instead of the interface identifier in the option 82 information.
See ["Including a Textual Description in DHCP Options" on page 1078](#).

Overriding Option 82 Information

You can configure the DHCP relay agent to add or remove the DHCP relay agent information option (option 82) in DHCP packets.

This feature causes the DHCP relay agent to perform one of the following actions, depending on the configuration:

- 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.

To override the default option 82 information in DHCP packets destined for a DHCP server:

1. Specify that you want to configure override options.

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

2. Specify that the option 82 information in DHCP packets is overwritten.

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

Including a Prefix in DHCP Options

When you configure the DHCP relay agent to include DHCP options in the packets that the relay agent sends to a DHCP server, you can specify that the relay agent add a prefix to the DHCP option. You can add a prefix to the following DHCP options:

- DHCPv4 option 82 Agent Circuit ID (suboption 1)
- DHCPv4 option 82 Agent Remote ID (suboption 2)
- DHCPv6 option 18 Relay Agent Interface-ID
- DHCPv6 option 37 Relay Agent Remote-ID

The prefix is separated from the DHCP option information by a colon (:), and it can include any combination of the `host-name`, `logical-system-name`, and `routing-instance-name` options. The DHCP relay agent obtains the values for the `host-name`, `logical-system-name`, and `routing-instance-name` as follows:

- If you include the `host-name` option, the DHCP relay agent uses the hostname of the device configured with the `host-name` statement at the `[edit system]` hierarchy level.
- If you include the `logical-system-name` option, the DHCP relay agent uses the logical system name configured with the `logical-system` statement at the `[edit logical-system]` hierarchy level.
- If you include the `routing-instance-name` option, the DHCP relay agent uses the routing instance name configured with the `routing-instance` statement at the `[edit routing-instances]` hierarchy level or at the `[edit logical-system logical-system-name routing-instances]` hierarchy level.

If you include the hostname and either or both of the logical system name and the routing instance name in the prefix, the hostname is followed by a forward slash (/). If you include both the logical system name and the routing instance name in the prefix, these values are separated by a semicolon (;).

The following examples show several possible formats for the DHCP option information when you specify the `prefix` statement for Fast Ethernet (fe) or Gigabit Ethernet (ge) interfaces with S-VLANs.

- If you include only the hostname in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
hostname:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

- If you include only the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

- If you include only the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

- If you include both the hostname and the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
host-name/logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

- If you include both the logical system name and the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

- If you include the hostname, logical system name, and routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

```
host-name/logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs but not S-VLANs, only the *vlan-id* value appears in the DHCP option format.

(DHCPv4) To configure a prefix with the option 82 information:

1. Specify that you want to configure option 82 support.

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

2. Configure DHCP relay agent to insert the Agent Circuit ID, the Agent Remote ID, or both.

- To configure the Agent Circuit ID:

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

- To configure the Agent Remote ID:

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

3. Specify that the prefix be included in the option 82 information. In this example, the prefix includes the hostname and logical system name.

- To include the prefix with the Agent Circuit ID:

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

- To include the prefix with the Agent Remote ID:

```
[edit forwarding-options dhcp-relay relay-option-82 remote-id]
user@host# set prefix host-name logical-system-name
```

(DHCPv6) To use a prefix with the DHCPv6 option 18 or option 37 information:

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

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

2. Configure DHCPv6 relay agent to insert option 18 (Relay Agent Interface-ID), option 37 (Relay Agent Remote-ID), or both.

- To configure option 18:

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

- To configure option 37:

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

3. Specify that the prefix is included in the option information. In this example, the prefix includes the hostname and logical system name

- To include the prefix with option 18:

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set prefix host-name logical-system-name
```

- To include the prefix with option 37:

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]
user@host# set prefix host-name logical-system-name
```

Including a Textual Description in DHCP Options

By default, when DHCP relay agent inserts option information in the packets sent to a DHCP server, the options include the interface identifier. However, you can configure the DHCP relay agent to include the textual description that is configured for the interface instead of the interface identifier. You can use the textual description for either the logical interface or the device interface.

You can include the textual interface description in the following DHCP options:

- DHCPv4 option 82 Agent Circuit ID (suboption 1)
- DHCPv4 option 82 Agent Remote ID (suboption 2)
- DHCPv6 option 18 Relay Agent Interface-ID
- DHCPv6 option 37 Relay Agent Remote-ID

The textual description is configured separately, using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. If you specify that the textual description is used and no description is configured for the interface, DHCP relay defaults to using the Layer 2 interface name.

In the case of integrated routing and bridging (IRB) interfaces, the textual description of the Layer 2 interface is used instead of the textual description of the IRB interface. If there is no description configured, the Layer 2 logical interface name is used.



NOTE: For IRB interfaces, the option 82 field must be able to uniquely identify the incoming interface based on either the Agent Circuit ID or Agent Remote ID . You can modify the information in the textual interface description to match the raw IFD

(physical interface without a subunit) name and configure the option 82 field to use the interface description.

You can use the textual description with the following DHCP options:

- DHCPv4 Option 82 Agent Circuit ID (suboption 1)
- DHCPv4 Option 82 Agent Remote ID (suboption 2)
- DHCPv6 Relay Agent Interface-ID (option 18)
- DHCPv6 Relay Agent Remote-ID (option 37)

(DHCPv4) To configure the DHCP relay option 82 suboption to include the textual interface description:

1. Specify that you want to configure option 82 support.

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

2. Configure DHCP relay agent to insert the Agent Circuit ID, Agent Remote ID, or both.

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

3. Specify that the textual description is included in the option 82 information. In this example, the option 82 information includes the description used for the device interface.

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

(DHCPv6) To configure the DHCPv6 option 18 or option 37 to include the textual interface description:

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

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

2. Configure DHCPv6 relay agent to insert option 18 (Relay Agent Interface-ID), option 37 (Relay Agent Remote-ID), or both.

- To configure option 18:

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

- To configure option 37:

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

3. Specify that the textual description is included in the option information. In the following example, the option information includes the description used for the device interface.

- To include the textual description in option 18:

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-interface-description device
```

- To include the textual description in option 37:

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]
user@host# set use-interface-description device
```

SEE ALSO

| [Configuring Interface Description](#)

Enable Processing of Untrusted Packets So Option 82 Information Can Be Used

By default, the DHCP relay agent treats client packets with a giaddr of 0 (zero) and option 82 information as if the packets originated at an untrusted source, and drops them without further processing. You can override this behavior and specify that the DHCP relay agent process DHCP client packets that have a giaddr of 0 (zero) and contain option 82 information.

To configure DHCP relay agent to trust option 82 information:

1. Specify that you want to configure override options.

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

2. Specify that the DHCP relay agent process DHCP client packets with a giaddr of 0 and that contain option 82 information.

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

Extracting an Option 82 or Option 37 Substring to Create an Interface Set

Starting in Junos OS Release 17.2R1, you can create an interface set based on a specific, delimited substring of the agent remote ID (ARI) string received in DHCP packets. Specify the predefined variable `$junos-pon-id-interface-set-name` in a dynamic profile to extract the substring from DHCPv4 (Option 82, suboption 2) or DHCPv6 (Option 37). This substring is inserted by the optical line terminal (OLT) in a passive optical network (PON) and is unique for that PON. The extracted substring is used as the name of the interface set.

The OLT must format the ARI string with a pipe symbol (|) as the delimiter between substrings. The substring extracted for the interface set name consists of the characters following the last delimiter in the ARI string. You determine the format and contents of the substring, and configure your OLT to insert the information. Typically, the substring may include the name and port of the OLT accessed by the CPE optical network terminal (ONT).

For example, the ARI format might be something like the following:

```
circuit-id|plan-name|ONT-serial-number|OLT-info
```

The following sample ARI strings follow that format:

```
ari-1001|100M|AAAA01234|ot101.xyz101-202
ari-9505|100M|AAAA01234|ot101.xyz101-202
ari-1238|100M|AAAA01234|ot101.xyz101-111
```

The first two ARIs share the same substring after the last delimiter, `ot101.xyz101-202`. The third ARI has a different last substring, `ot101.xyz101-111`. The predefined variable extracts both of these substrings. Two interface sets are created, named `ot101.xyz101-202` and `ot101.xyz101-111`.

The two customer circuits identified by ot101.xyz101-202 are aggregated into that interface set. The single circuit identified by ot101.xyz101-111 is associated with the other set. The interface sets can subsequently be used to apply CoS and services to their associated subscriber circuits.

Before you begin:

- Configure your OLTs to provide an agent remote ID string in the required format.
- Configure your DHCPv4 or DHCPv6 relay agents to insert the agent remote ID received from the OLT for forwarding to the DHCP local server.
- Create the dynamic profile.

This procedure shows only the configuration required for specifying the predefined variable.

1. Access the desired dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Specify the predefined variable to create the interface set.

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

3. Complete the dynamic profile configuration.

You can use the show subscribers extensive command to display the interface set name and the complete ARI string.

show subscribers extensive (Passive Optical Network Circuit Interface Set)

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
...
Interface Set: ot101.xyz101-202
...
Agent Remote ID: ABCD01234|100M|AAAA01234|ot101.xyz101-202
...
```

SEE ALSO

Junos OS Predefined Variables
Configuring Predefined Dynamic Variables in Dynamic Profiles
Configuring a Basic Dynamic Profile

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.2R1	Starting in Junos OS Release 17.2R1, you can create an interface set based on a specific, delimited substring of the agent remote ID (ARI) string received in DHCP packets.

RELATED DOCUMENTATION

DHCP Overview 1015
DHCPv6 Local Server 1449
DHCPv6 Relay Agent 1455
Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings 1033

Default Services for DHCP Subscribers

IN THIS SECTION

- [Default Subscriber Service Overview | 1091](#)
- [Configuring a Default Subscriber Service | 1092](#)

Default Subscriber Service Overview

Subscriber management enables you to specify a default subscriber service for DHCP subscribers. The default service (dynamic profile) is applied to subscribers when the subscriber logs in. By configuring a default service, you can apply a particular service (for example, a basic service) to subscribers who are not explicitly assigned a service.

When a subscriber logs in, the configured default service is always activated, even when remote service provisioning or RADIUS service activation is configured for the subscriber. The default service is deactivated only when the subscriber is successfully provisioned by the PCRF by means of the GX-Plus application. (Remote provisioning is configured by the provisioning-order statement at the [edit access profile] hierarchy level.)

In all other cases, the default service remains active. For example, if RADIUS authentication is configured but service activation is not, the default subscriber service remains activated. Likewise, if RADIUS authentication is not configured, the default subscriber service remains activated.

Default services can also be deactivated either with a RADIUS CoA deactivate request or with the request network-access aaa subscriber delete session-id command.

To create and assign a default subscriber service, you must complete the following operations:

- Create the service—Ensure that the service you want to use has been configured in a dynamic service profile. The actual service is no different than any other service used for subscriber management.
- Specify the default service—Use the Junos OS CLI to specify the service that is used as the default service.
- Specify the interfaces on which the default service is assigned —Use the Junos OS CLI to specify that the default service is used globally, for a group of interfaces, or for a specific interface.

Configuring a Default Subscriber Service

Subscriber management enables you to specify a default subscriber service for DHCP (and DHCPv6) local server and DHCP relay agent. The default service is the service (dynamic profile) that is applied to subscribers when they log in.

Default services are subsequently deactivated in any of the following circumstances:

- A PCRF responds to AAA for the subscriber.
- A RADIUS CoA deactivation request is issued.
- You deactivate the service manually through the CLI.

To configure a default subscriber service:

1. Ensure that the service you want to use as the default has been configured in a dynamic profile.
2. Specify the default service.

The following example configures the default service for DHCP local server subscribers.

```
[edit system services dhcp-local-server]
user@host# set service-profile retailer1-subscriber
```

3. Attach the default service—you can attach the profile globally, for a group of interfaces, or for a specific interface.

The following example attaches the profile to a named group of interfaces for DHCP local server.

- Specify the group to which the default service is attached.

```
[edit system services dhcp-local-server]
user@host# set group subscriber-svl
```

- Specify the dynamic profile that defines the default service.

```
[edit system services dhcp-local-server group subscriber-svl]
user@host# set service-profile retailer1-subscriber
```

SEE ALSO

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

RELATED DOCUMENTATION

Service Activation and Deactivation Using the CLI Instead of RADIUS

[Gx-Plus for Provisioning Subscribers | 1781](#)

DHCP Client Attribute and Address Assignment

IN THIS SECTION

- [DHCP Attributes Overview | 1094](#)
- [Attributes That Can Be Applied to DHCP Clients | 1095](#)
- [Configuring DHCP Attributes for All Clients or a Group of Clients | 1098](#)
- [Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 1100](#)
- [Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | 1101](#)

- [Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option | 1103](#)
- [Specifying the Subnet for DHCP Client Address Assignment | 1104](#)
- [DHCP Local Server Handling of Client Information Request Messages | 1104](#)
- [Enabling Processing of Client Information Requests | 1105](#)
- [DNS Address Assignment Precedence | 1106](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching | 1106](#)

DHCP Attributes Overview

IN THIS SECTION

- [Benefits of Configuring DHCP Attributes | 1095](#)

You can configure features that are specific to the DHCP application that are applied to only certain DHCP clients or to all DHCP clients with DHCP attributes. DHCP uses the attributes to determine the scope of the client operation. For example, you can configure attributes that set the maximum lease time or preferred lifetime of the lease, the domain in which to search for DHCP servers, match criteria for which address range to use from within an address pool, and so on. You might specify that a client that matches certain prerequisite information is dynamically assigned an address from a particular named address range. Based on which named range is used, DHCP specifies additional DHCP attributes.

You can configure DHCP attributes in the following ways:

- On the RADIUS server so that they are conveyed in the corresponding DHCP option when a subscriber is authenticated. Refer to your RADIUS server documentation for more information.
- For specific DHCPv4 or DHCPv6 clients that receive an address from the local address assignment pool with the `dhcp-attributes` statement at the `[edit access address-assignment pool pool-name]` hierarchy level.
- As a set of attributes that you can apply to DHCP clients outside of specific address pools. Define the attribute set with the `protocol-attributes` statement at the `[edit access]` hierarchy level. Then apply the set with a different `protocol-attributes` statement to any of the following:
 - For all DHCPv4 clients at the `[edit system services dhcp-local-server overrides]` hierarchy level.

- For a group of DHCPv4 clients at the [edit system services dhcp-local-server group *group-name* overrides] hierarchy level.
- For all DHCPv6 clients at the [edit system services dhcp-local-server dhcpv6 overrides] hierarchy level.
- For a group of DHCPv6 clients at the [edit system services dhcp-local-server dhcpv6 group *group-name* overrides] hierarchy level.

The DHCP local server processes attributes provided by different methods in the following hierarchy:

RADIUS > address pool > global > other

1. When the attribute is configured in RADIUS, the value in the corresponding option received by the DHCP local server is used.
2. When the attribute is configured for an address pool, that value is used for clients assigned addresses from that pool.
3. When the attribute is configured globally with the protocol-attributes statement, that value is used for all clients.
4. When none of the other criteria is met but the attribute is configured at the [edit access] hierarchy level, that value is used for all clients. If the attribute is configured at the [edit access profile] hierarchy level, that value is used for clients using the profile.

Benefits of Configuring DHCP Attributes

You can match desired attributes to specific clients based on matching criteria. You have the flexibility to assign attributes and values when an address is assigned from a pool, globally for clients not using address pools, or with RADIUS attributes at authentication.

Attributes That Can Be Applied to DHCP Clients

This topic provides descriptions of DHCPv4 and DHCPv6 options.

[Table 33 on page 1095](#) describes the DHCPv4 client attributes that you can configure.

Table 33: DHCP Attributes

Attribute	Description	DHCP Option
boot-file	Boot filename advertised to the client, and used by the client to complete configuration.	67

Table 33: DHCP Attributes *(Continued)*

Attribute	Description	DHCP Option
boot-server	Boot server containing the boot file.	66
domain-name	Domain in which clients search for a DHCP server host.	15
grace-period	Grace period offered with the lease.	-
maximum-lease-time	Maximum lease time allowed by the DHCP server.	51
name-server	IP address of DNS server to which clients can send DNS queries.	6
netbios-node-type	NetBIOS node type.	46
option	User-defined options.	-
option-match	Option 82 value is mapped to named address range.	-
router	IP address for routers on the subnetwork.	3
server-identifier	IP address used as the DHCP source address	54
t1-percentage	Percentage of the preferred-lifetime that the client (router) waits before sending renew messages to the DHCPv4 server that granted the original lease to extend the client's lease.	58
t1-renewal-time	Time that the client (router) waits before sending renew messages to extend the client's lease. The renew messages are sent to the DHCPv4 server that granted the original lease. This attribute is an alternative to t1-percentage.	58

Table 33: DHCP Attributes (Continued)

Attribute	Description	DHCP Option
t2-percentage	Percentage of the preferred-lifetime that the client (router) waits before sending rebind messages to any available DHCPv4 server to extend the client's lease.	59
t2-rebinding-time	Time that the client (router) waits before sending rebind messages to extend the client's lease. The rebind messages are sent to any available DHCPv4 server. This attribute is an alternative to t2-percentage.	59
tftp-server	Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file.	150
wins-server	IP address of the Windows NetBIOS name server.	44

[Table 34 on page 1097](#) describes the DHCPv6 client attributes that you can configure.

Table 34: DHCPv6 Attributes

Attribute	Description	DHCPv6 Option
dns-server	IPv6 address of DNS server to which clients can send DNS queries.	23
grace-period	Grace period offered with the lease.	–
maximum-lease-time	Maximum lease time allowed by the DHCP server.	–
option	User-defined options.	–
preferred-lifetime	Length of time that a valid address is in the preferred state. When the preferred lifetime expires, the address becomes deprecated.	–

Table 34: DHCPv6 Attributes (Continued)

Attribute	Description	DHCPv6 Option
sip-server-address	IPv6 address of SIP outbound proxy server.	22
sip-server-domain-name	Domain name of the SIP outbound proxy server.	21
t1-percentage	Percentage of the preferred-lifetime that the client (router) waits before sending renew messages to the DHCPv6 server that granted the original lease to extend the client's lease.	-
t1-renewal-time	Time that the client (router) waits before sending renew messages to extend the client's lease. The renew messages are sent to the DHCPv6 server that granted the original lease. This attribute is an alternative to t1-percentage.	-
t2-percentage	Percentage of the preferred-lifetime that the client (router) waits before sending rebind messages to any available DHCPv6 server to extend the client's lease.	-
t2-rebinding-time	Time that the client (router) waits before sending rebind messages to extend the client's lease. The rebind messages are sent to any available DHCPv6 server. This attribute is an alternative to t2-percentage.	-
valid-lifetime	Length of time that the address remains in the valid state. When the lifetime expires, the address becomes invalid.	-

Configuring DHCP Attributes for All Clients or a Group of Clients

You can configure DHCP client attributes to determine the scope of the client operation. For example, you can configure attributes that set the maximum lease time or preferred lifetime of the lease, the domain in which to search for DHCP servers, the match criteria that determine the address range to use from within an address pool, and so on.

You can configure DHCP attributes to be applied to clients in the following ways:

- Globally to all clients or only to clients in specific groups.
- By an address-assignment pool; the attributes apply only to clients that receive addresses from a specific address assignment pool. See "[Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address](#)" on page 1100 for more information about this method.

To assign attributes globally or to a group:

1. Create a DHCP attribute set that you want to apply to clients with the DHCPv4 or DHCPv6 local server.

```
[edit access]
user@host# edit protocol-attributes attribute-set-name
```

2. Specify the attributes to include in the attribute set.

```
[edit access protocol-attributes attribute-set-name]
user@host# set attribute
```

3. Apply the attribute set to the desired DHCP clients.

- To all DHCPv4 clients:

```
[edit system services dhcp-local-server overrides]
user@host# set protocol-attributes attribute-set-name
```

- To a group of DHCPv4 clients:

```
[edit system services dhcp-local-server group group-name overrides]
user@host# set protocol-attributes attribute-set-name
```

- To all DHCPv6 clients:

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set protocol-attributes attribute-set-name
```

- To a group of DHCPv6 clients:

```
[edit system services dhcp-local-server dhcpv6 group group-name overrides]
user@host# set protocol-attributes attribute-set-name
```

For example, the following configuration creates an attribute set named `attr-v4-1` and applies the set to all DHCPv4 clients.

```
[edit]
user@host# set access protocol-attributes attr-v4-1 maximum-lease-time seconds
user@host# set access protocol-attributes attr-v4-1 t1-renewal-time 120000
user@host# set system services dhcp-local-server overrides protocol-attributes attr-v4-1
```

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address

You use the address-assignment pool feature to include DHCP attributes specific to the client when clients obtain an address. The DHCP client application uses the attributes to determine how addresses are assigned, and to also provide optional characteristics to the client. For example, the DHCP application might specify that a client that matches certain prerequisite information is dynamically assigned an address from a particular named range. Based on which named range is used, DHCP specifies additional DHCP attributes such as the boot file that the client uses, the lease grace period, and the maximum lease time.

You use the `dhcp-attributes` statement to configure DHCP client-specific attributes for address-assignment pools. ["Attributes That Can Be Applied to DHCP Clients" on page 1095](#) describes the supported attributes you can configure for IPv4 and IPv6 address-assignment pools (or optionally assign to all clients or clients in a group).

To configure address-assignment pool attributes for DHCP clients:

1. Specify the name and IP family of the address-assignment pool.

```
[edit access]
user@host# edit address-assignment pool pool-name family inet
```

2. Configure optional DHCP client attributes.

```
[edit access address-assignment pool isp_1 family inet]
user@host# set dhcp-attributes attribute1 value1 attribute 2 value2 attribute 3 value3
```

For example, the following configuration specifies values for the boot server, grace period, and maximum lease time for the isp1 pool for DHCPv4:

```
[edit access]
user@host# edit address-assignment pool isp_1 family inet
[edit access address-assignment pool isp_1 family inet]
user@host# set dhcp-attributes boot-server 192.168.200.100 grace-period 3600 maximum-lease-time
18000
```



NOTE: The DNS name server addresses that are configurable as DHCP attributes can also be configured globally at the routing instance level and in access profiles. For more information, see ["DNS Name Server Address Overview" on page 1503](#).

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

You can specify the match order in which the extended DHCP local server uses the client data to determine the address-assignment pool that provides the IP address and configuration for a DHCP client. If you do not specify any pool match order, the device uses the default IP address configured in IP address first matching option to select the address pool.

Example:

```
[edit system services dhcp-local-server]
user@host# set pool-match-order
```

You can specify the order for pool matching methods. You can specify the methods in any order. All methods are optional. IP address first method is default method.

- **IP address first**—Default option. The server selects the address-assignment pool to use by matching the IP address in the client DHCP request with the network address of the address-assignment pool.
- If the client request contains the gateway IP address (giaddr), the local server matches the giaddr to the address-assignment pool's address.
- If the client request does not contain the giaddr, then the DHCP local server matches the IP address of the receiving interface to the address of the address-assignment pool.

Example:

```
[edit system services dhcp-local-server pool-match-order]
user@host# set ip-address-first
```

- External authority—The DHCP local server receives the address assignment from an external authority, such as RADIUS or Diameter.
- If RADIUS is the external authority, the DHCP local server uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool.
- If Diameter is the external authority, the server uses the Diameter counterpart of the Framed-IPv6-Pool attribute to determine the pool.

Example:

```
[edit system services dhcp-local-server pool-match-order]
user@host# set external-authority
```

- Option 82—For IPv4 address-Extended DHCP local server matches the DHCP relay agent information option (option 82) in the client DHCP packets to a named range in the address-assignment pool. Named ranges are subsets within the overall address-assignment pool address range, which you can configure when you create the address-assignment pool.

Example:

```
[edit system services dhcp-local-server pool-match-order]
user@host# set option-82
```

To use the DHCP local server option 82 matching feature with an IPv4 address-assignment pool, you must ensure that the `option-82` statement is included in the `dhcp-attributes` statement for the address-assignment pool.

This example shows an extended DHCP local server configuration that includes optional IPv4 address-assignment pool matching and interface groups. For pool matching, this configuration specifies that the DHCP local server first check the response from an external authentication authority (for example, RADIUS) and use the Framed-IPv6-Pool attribute to determine the address-assignment pool to use for the client address. If no external authority match is found, the DHCP local server then uses `ip-address-first` matching together with the option 82 information to match the named address range for client IPv4

address assignment. The option 82 matching must also be included in the address-assignment pool configuration.

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
    interface fe-0/0/2.1;
  }
  group group_two {
    interface fe-0/0/3.0;
    interface fe-0/0/3.1;
  }
  pool-match-order {
    external-authority
    ip-address-first;
    option-82;
  }
}
```

Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option

Subscriber management or DHCP management enables you to specify that DHCP local server assign a particular address to a client. For example, if a client is disconnected, you might use this capability to assign the same address that the client was using prior to being disconnected. If the requested address is available, DHCP assigns it to the client. If the address is unavailable, the DHCP local server offers another address, based on the address allocation process.

Both DHCP local server and DHCPv6 local server support the specific address request feature. DHCP local server uses DHCP option 50 in DHCP discover messages to request a particular address, while DHCPv6 local server uses the IA_NA option (Identity Association for Non-Temporary Addresses) in DHCPv6 solicit messages.



NOTE: Subscriber management (DHCP management) supports only one address for each of the DHCPv6 IA_NA or IA_PD address types. If the DHCPv6 client requests more than one address for a given type, the DHCPv6 local server uses only the first address and ignores the other addresses.

Specifying the Subnet for DHCP Client Address Assignment

Subscriber management enables you to explicitly specify the subnet to which the DHCP local server matches the requested IP address. The server accepts and uses an active client's requested IP address for address assignment only when the requested address and the IP address of the DHCP server interface are in the same subnet. The server accepts and uses a passive client's requested IP address only when the requested address and the IP address of the relay interface are in the same subnet. The DHCPv6 local server supports the same process for DHCPv6 clients and addresses.

To specify the subnet used for client address assignment:

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set requested-ip-network-match 10
```

- For DHCPv6 local server:

```
[edit forwarding-options dhcp-local-server dhcpv6]
user@host# set requested-ip-network-match 30
```

DHCP Local Server Handling of Client Information Request Messages

DHCP clients that already have externally provided addresses may solicit further configuration information from a DHCP server by sending a DHCP inform or DHCPv6 information-request message that indicates what information is desired. These message types can be collectively referred to as information request messages. By default, DHCP local server and DHCPv6 local server ignore any DHCP information requests that they receive. You can override this default behavior to enable processing of these messages.

If you enable processing of information requests, DHCP local server responds to the client with a DHCP acknowledgment message that includes the requested information—if it is available. DHCPv6 local server responds in the same manner but uses a DHCP reply message. No subscriber management or DHCP-management is applied as a result of the DHCP information request message.

By default, DHCP relay and DHCP relay proxy automatically forward DHCP information request messages without modification if the messages are received on an interface configured for a DHCP server group. DHCP relay and relay proxy drop information request messages received on any other interfaces. You cannot disable this default DHCP relay and relay proxy behavior.

The information requested by these clients is typically configured with the `dhcp-attributes` statement for an address pool defined by the address-assignment `pool pool-name` statement at the `[edit access]` hierarchy level.

When you enable processing of DHCP information requests, you can optionally specify the name of the pool from which the local server retrieves the requested configuration information for the client. If you do not specify a local pool, then the local server requests that AAA selects and returns only the name of the relevant pool.



NOTE: PPP interfaces are not supported on EX Series switches.

When DHCPv6 is configured over PPP interfaces, the PPP RADIUS authentication data can be used to select the pool from which the response information is taken. Additionally other RADIUS attributes can also be inserted into the DHCPv6 reply message. If an overlap exists between RADIUS attributes and local pool attributes, the RADIUS values are used instead of the local configuration data. If no RADIUS information is received from the underlying PPP interface, then the behavior is the same as described previously for non-PPP interfaces.

Enabling Processing of Client Information Requests

Configure one or more local address pools if you want to use a local pool rather than one provided by AAA. See [DHCPv6 Address-Assignment Pools](#). For processing information request messages, the address configuration is not necessary. For DHCP local server, you must specify the IPv4 family; for DHCPv6 local server, you must specify the IPv6 family.

See ["Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address" on page 1100](#) for details about how to configure the information sought by clients that send information request messages.

By default, DHCP local server and DHCPv6 local server do not respond to information request (DHCP inform and DHCPv6 information-request) messages from the client. You can enable DHCP local server and DHCPv6 local server to process these messages and respond to them with an acknowledgment (ack or reply message, respectively) and the requested information.

DHCP relay agent automatically forwards the information request messages without modification to the configured server group by means of the interfaces configured for the respective server group. The messages are dropped if they are received on an unconfigured interface. DHCP relay proxy also supports forwarding these messages. You cannot disable forwarding of the information request messages.

To enable processing of DHCP client information request messages:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server overrides]
user@host# set process-inform
```


- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set process-inform
```

2. (Optional) Specify a pool name from which DHCP information is returned to the client.

- For DHCP local server:

```
[edit system services dhcp-local-server overrides process-inform]
user@host# set pool pool-name
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6 overrides process-inform]
user@host# set pool pool-name
```

DNS Address Assignment Precedence

Subscriber management supports four methods for assigning addresses to DHCP clients. When multiple methods are configured, the router uses the following precedence order to determine which address to assign to the client.

1. Address defined on the RADIUS server by Internet Assigned Numbers Authority (IANA) vendor ID 4874 attributes 26-4 (Primary-DNS) and 26-5 (Secondary-DNS).
2. Address defined on the RADIUS server by IANA vendor ID 2636 attributes 26-31 (Primary-DNS) and 26-33 (Secondary-DNS).
3. Address defined on the RADIUS server by IANA vendor ID 311 attributes 26-28 (MS-Primary-DNS-Server) and 26-29 (MS-Secondary-DNS-Server).
4. Address defined in the local address pool on the router.

Example: Extended DHCP Local Server Configuration with Optional Pool Matching

This example shows an extended DHCP local server configuration that includes optional IPv4 address-assignment pool matching and interface groups. For pool matching, this configuration specifies that the DHCP local server first check the response from an external authentication authority (for example, RADIUS) and use the Framed-IPv6-Pool attribute to determine the address-assignment pool to use for the client address. If no external authority match is found, the DHCP local server then uses ip-address-first matching together with the option 82 information to match the named address range for client IPv4

address assignment. The option 82 matching must also be included in the address-assignment pool configuration.

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
    interface fe-0/0/2.1;
  }
  group group_two {
    interface fe-0/0/3.0;
    interface fe-0/0/3.1;
  }
  pool-match-order {
    external-authority
    ip-address-first;
    option-82;
  }
}
```



NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Address-Assignment Pools for Subscriber Management | 1476](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[Standard and Vendor-Specific RADIUS Attributes | 3](#)

DHCP Lease Times for IP Addresses

IN THIS SECTION

- [DHCP Lease Timers | 1108](#)
- [DHCP Lease-Time Validation Overview | 1109](#)
- [Configuring a DHCP Lease-Time Threshold | 1111](#)
- [DHCP Asymmetric Leasing Overview | 1112](#)
- [Configuring DHCP Asymmetric Leasing | 1113](#)

DHCP Lease Timers

IN THIS SECTION

- [Benefits of Configuring DHCP Timers in Address Pools | 1109](#)

Subscriber management supports configurable timers that you can use to manage the DHCPv4 and DHCPv6 address leases provided by address-assignment pools. In addition to the maximum-lease-time timer, which sets the maximum time for which the DHCP local server can grant a lease, you can use DHCP client-specific attributes to configure timers that govern the lifetimes of existing leases that have been obtained from an address-assignment pool. Starting in Junos OS Release 17.2R1, this feature is supported for both DHCPv4 and DHCPv6; in earlier releases, only DHCPv6 is supported.

The following list describes the configurable timers for DHCPv4 and DHCPv6 address-assignment pools:

- **preferred-lifetime**—Length of time that a valid address is in the preferred state and can be used without any restrictions. When the preferred-lifetime expires, the address becomes deprecated. A deprecated address should not be used for new communications, but might continue to be used for existing communications in certain cases.

If the valid-lifetime is also configured, the preferred-lifetime must be less than the valid-lifetime. The preferred-lifetime and the maximum-lease-time are mutually exclusive and cannot both be configured.

- **valid-lifetime**—Length of time that an address remains in the valid state, during which the address can be used for new or existing communications. When the valid-lifetime expires, the address becomes invalid, and can no longer be used.

If the preferred-lifetime is also configured, the valid-lifetime must be greater than the preferred-lifetime. The valid-lifetime and the maximum-lease-time are mutually exclusive and cannot both be configured.

- **t1 percentage**—Percentage of the preferred-lifetime that the client waits before contacting the DHCP local server that originally granted the lease to request that the address lease be extended. T1 is also called the renewal time.
- **t1-renewal-time**—Time in seconds that the client waits before contacting the DHCP local server that originally granted the lease to request that the address lease be extended. T1 is also called the renewal time.
- **t2 percentage**—Percentage of the preferred-lifetime that the client waits before sending a request to any available DHCP local server to extend the address lease. T2 is also called the rebind time.
- **t2-rebinding-time**—Time in seconds that the client waits before broadcasting a request to all available DHCP local servers to request that the address lease be extended. T2 is also called the rebind time.

Benefits of Configuring DHCP Timers in Address Pools

Using address pools to configure values for these timers gives you fine-grained control over which clients get specific values.

DHCP Lease-Time Validation Overview

IN THIS SECTION

- [Benefits of DHCP Lease Time Validation | 1111](#)

In a subscriber access environment, a DHCP server obtains an address lease from either local configuration or from an external DHCP server, and assigns the lease to the DHCP client address.

Obtaining leases from external sources can present issues when the external source is owned or managed by a third party—the third party might configure the external source to provide address leases that are unsuitable for the subscriber access environment. For example, extremely short lease times can create unnecessary traffic that results in reduced performance in the network.

To avoid potential issues caused by short DHCP lease times, subscriber management provides a lease-time validation feature. Lease-time validation enables you to explicitly configure a threshold for the minimum lease time allowed in your subscriber access environment, and to specify a violation action (such as dropping the lease offer) the router takes when a short lease time is offered by a third party. You can specify the following violation actions:

- **drop**—(DHCPv4 and DHCPv6 relay agent) The third-party lease offer is dropped and the client binding fails.
- **override-lease**—(DHCPv4 and DHCPv6 local server) The third-party lease time is overridden by the specified threshold value.
- **strict**—(DHCPv4 and DHCPv6 local server) The third-party lease is ignored and the client binding fails.
- **no action**—If you do not specify a violation action, DHCP binds the client using the third-party lease but marks the binding as lease-time violating.

A lease-time violation can occur during the initial lease grant or during a rebinding or renewal operation. To reduce excessive and redundant log messages, the router consolidates lease-time violation reporting, as shown in [Table 35 on page 1110](#).

Table 35: Lease-Time Violation Event Logging

Event	syslog	Extended DHCP Traceoptions
Initial lease-time violation for the specific DHCP server	warning	warning
Number of lease-time violations return to zero for the specific DHCP server	warning	warning
Status of lease-time violations caused by specific DHCP server, reported in the interval configured in <code>ltv-syslog-interval</code> command	warning	–
Violation action of drop occurred, or the DHCP packet was not generated	–	warning
Violation action of override-lease occurred (DHCP local server only)	–	warning
Lease-time violation	–	warning

Benefits of DHCP Lease Time Validation

- Enables you to avoid unnecessary traffic that can reduce performance when third-party DHCP leases are too short by configuring a minimum allowed lease time and actions taken for invalid leases.

Configuring a DHCP Lease-Time Threshold

Starting in Junos OS Release 14.1, subscriber management provides a lease-time validation feature that enables you to specify the minimum DHCP lease time allowed in your subscriber access environment. When you configure lease-time validation, you specify the lease-time threshold and the action the router performs when an offered lease time is less than the threshold (such as dropping the lease).

Lease-time validation ensures that leases that are offered by third-party DHCP servers or address assignment pools always meet the requirements of your network. For example, you want short leases to be rejected because they can result in excessive renewal traffic that can impact network performance.

You can configure lease-time validation on DHCPv4 and DHCPv6 local servers, and DHCPv4 and DHCPv6 relay agents, and for individual interfaces or interface groups. DHCP relay proxy also supports lease-time validation.

The following procedure describes the steps you use to configure lease-time validation. This example describes a configuration for DHCP relay agent. You use similar steps at the appropriate hierarchy levels for DHCP local server, DHCPv6 local server, and DHCPv6 relay agent.

To configure lease time validation for DHCP relay agent:

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

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

2. Specify that you want to configure the DHCP lease time validation feature.

```
[edit forwarding-options dhcp-relay]
user@host# set lease-time-validation
```

3. Configure the threshold that specifies the minimum DHCP client lease time allowed in your network.

```
[edit forwarding-options dhcp-relay lease-time-validation]
user@host# set lease-time-threshold 3600
```

4. Configure the action the router takes when a lease time violation occurs.

```
[edit forwarding-options dhcp-relay lease-time-validation]
user@host# set violation-action drop
```



NOTE: DHCP relay agent and DHCP local server support different violation actions. See the `violation-action` statement for descriptions of the actions.

If a lease violation occurs when you have not configured a violation action, DHCP binds the client using the third-party lease. DHCP then marks the binding as having violated the lease time.

5. (Optional) Configure how often you want the router to consolidate and log syslog warning messages.

```
[edit system processes dhcp-service]
user@host# set ltv-syslog-interval 3600
```

DHCP Asymmetric Leasing Overview

IN THIS SECTION

- [Benefits of Asymmetric DHCP Lease Timing | 1113](#)

Starting in Junos OS Release 17.1R1, *asymmetric leasing* provides a way to send a DHCP client a lease that is shorter than the actual lease granted by the DHCP local server. In some networks, you might need to change an existing DHCP address assignment before the granted lease time has expired, or learn as soon as possible that a client is no longer using an address. The shorter lease is called the *asymmetric lease* or the *short lease*. The originally granted lease is called the *long lease* or simply, the lease.

You can configure asymmetric leasing on either the DHCP relay agent or the DHCP local server. In a typical configuration, you configure it on the relay agent. When the DHCP local server receives a discover packet from a DHCP client, it returns an offer packet to the client. The client selects a local server and requests an address assignment. The DHCP local server sends an acknowledgment packet containing the address, lease duration and other information to the client. Rather than forwarding this packet, the relay agent saves the lease information and then generates a new acknowledgment packet with a short lease and forwards that to the client.

When the DHCP client makes a subsequent request for lease renewal, the relay agent does not pass the request to the local server. Instead, the relay agent recreates the short lease from the saved information and returns it to the client in an acknowledgment packet. The relay agent continues to renew the short leases for the client until the long lease renew time expires. By default, the long lease renew time is equal to one-half the duration of the long lease.

When the long lease renew time expires, the asymmetric lease is no longer valid. Subsequent renewal requests from the client are forwarded by the relay agent to the local server. If the local server acknowledges the request, it renews the long lease and the process begins again, with the relay agent generating a short lease for the client instead of sending the long lease. Otherwise, the lease is not renewed.

With asymmetric leasing, there is also a renew time for the short lease. The client sends renewal requests to the relay agent at intervals equal to the short lease renew time. By default, this period is equal to one-half the short lease duration. If the DHCP client does not request renewal of the lease before the short lease time expires, the relay agent notifies the local server that the lease is no longer in use and the address can be reassigned.

Benefits of Asymmetric DHCP Lease Timing

- Allows the early renewal of DHCP client leases without requiring device support for a forced renewal. Because the lease can be administratively rescinded on the DHCP local server or DHCP relay agent, the next short-duration lease refresh cycle can trigger a full renegotiation of the client lease. This capability reduces the number of devices that must be managed to trigger a new address allocation cycle.
- Enables early detection of inactive client leases. The asymmetric lease includes an upstream notification from the DHCP client back to the lease grantor, effectively providing a liveness detection that enables unused addresses to be reclaimed sooner. Because this mechanism is all within DHCP, the service provider does not have to rely on the set of devices involved in address allocation to be configured with other protocols that support some kind of liveness detection, such as bidirectional forwarding detection (BFD).

Configuring DHCP Asymmetric Leasing

You can configure asymmetric leasing to provide a DHCP or DHCPv6 client with a lease that is shorter than the lease granted by the DHCP local server. The shorter lease duration means that the client must renew the lease more frequently. When it does not renew the lease, the address is freed up sooner than when the client uses the original long lease. You configure asymmetric leasing by overriding the DHCP configuration at the global level, for a named group of interfaces, or for a specific interface within a named group.



NOTE: For simplicity, the procedures in this topic show only the global-level configuration. For information about overriding the DHCP configuration at other levels, see ["Overriding the Default DHCP Relay Configuration Settings" on page 1036](#) and ["Overriding the Default DHCP Local Server Configuration Settings" on page 1034](#).

You can configure asymmetric leasing on either the DHCP relay agent or the DHCP local server. For most networks, the relay agent configuration is more useful.

To configure asymmetric leasing for DHCP relay agent for DHCPv4:

1. Specify that you want to configure override options.

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

2. Configure the duration of the short (asymmetric) lease.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set asymmetric-lease-time seconds
```

To configure asymmetric leasing for DHCP relay agent for DHCPv6:

1. Specify that you want to configure override options.

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

2. Configure the duration of the short (asymmetric) lease for DHCPv4 clients and separately for DHCPv6 clients.

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set asymmetric-lease-time seconds
user@host# set asymmetric-prefix-lease-time seconds
```

To configure asymmetric leasing for DHCP local server for DHCPv4:

1. Specify that you want to configure override options.

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

2. Configure the duration of the short (asymmetric) lease.

```
[edit system services dhcp-local-server overrides]
user@host# set asymmetric-lease-time seconds
```

To configure asymmetric leasing for DHCP local server for DHCPv6:

1. Specify that you want to configure override options.

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

2. Configure the duration of the short (asymmetric) lease for DHCPv4 clients and separately for DHCPv6 clients.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set asymmetric-lease-time seconds
user@host# set asymmetric-prefix-lease-time seconds
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.2R1	Starting in Junos OS Release 17.2R1, this feature is supported for both DHCPv4 and DHCPv6; in earlier releases, only DHCPv6 is supported.
17.1	Starting in Junos OS Release 17.1R1, <i>asymmetric leasing</i> provides a way to send a DHCP client a lease that is shorter than the actual lease granted by the DHCP local server.
14.1	Starting in Junos OS Release 14.1, subscriber management provides a lease-time validation feature that enables you to specify the minimum DHCP lease time allowed in your subscriber access environment.

RELATED DOCUMENTATION

[DHCP Client Attribute and Address Assignment | 1093](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

DHCP Leasequery Methods

IN THIS SECTION

- [Benefits of DHCP Leasequery | 1117](#)
- [DHCP Individual Leasequery | 1117](#)
- [DHCP Bulk Leasequery | 1121](#)
- [DHCP Active Leasequery | 1127](#)
- [Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations | 1138](#)
- [Configuring and Using DHCP Individual Leasequery | 1139](#)
- [Configuring and Using DHCP Bulk Leasequery | 1141](#)
- [Configuring and Using DHCP Active Leasequery | 1145](#)
- [Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | 1152](#)
- [Verifying and Managing DHCP Individual and Bulk Leasequery Configurations | 1157](#)
- [Verifying and Managing DHCP Active Leasequery Operations | 1158](#)

In a subscriber access network, a DHCP local server maintains a significant amount of binding information related to the IP addresses or DHCPv6 delegated prefixes that the server has leased to DHCP clients. When DHCP clients are connected to the DHCP server by way of a DHCP relay agent, the DHCP relay agent gleans data from the DHCP packets it forwards, such as IP address, necessary to reach the endpoint. The relay agent maintains lease and route information relevant to the DHCP clients. The relay agent uses that information when providing subscriber services for the clients. When the relay agent is restarted or when the agent host device is rebooted or replaced, the relay agent loses that information. You can use a request command to trigger the relay agent to send a leasequery message to the local server to recover the binding information for DHCP clients so that the relay agent can restore its lease information database.

Subscriber management supports the following types of leasequery operations:

- Individual leasequery—Provides lease information for a single binding on request (query and response mode).
- Bulk leasequery—Provides lease information for multiple bindings on request (query and response mode).
- Active leasequery—Provides a stream of live updates for multiple bindings when configured.

Benefits of DHCP Leasequery

- Leasequery provides a lightweight way for a DHCPv4 or DHCPv6 relay agent to recover the authoritative location information related to leased DHCP IP/IPv6 addresses and delegated prefixes from the DHCP local server when the relay agent has been restarted or replaced.
- Bulk leasequery removes the need to query individual bindings for specific clients, allowing a single request to return information for hundreds or thousands of subscribers. This method does not wait for data traffic to trigger a query, so it scales better than individual leasequery when the agent has thousands of clients. In the case of DHCPv6, the relay agent may not be able to form individual queries.
- Active leasequery provides continual live updates of binding information to one or more relay agents when configured. In addition to updates between relay agent and local server, you can configure a peering relationship between relay agents. This enables the peers to continually synchronize their binding information with each other, providing redundancy if a peer goes down or is rebooted. The active peer immediately maintains service for the clients that were using the affected relay agent.
- Topology discovery enables relay agent peers on BNGs configured for M:N subscriber redundancy to automatically build translation tables so that subscriber redundancy groups continue to be served when a primary BNG fails over to a backup. This automatic behavior frees you from having to build tables statically. Static configuration is error-prone in scaled networks and does not adapt dynamically to changes in the network.

DHCP Individual Leasequery

IN THIS SECTION

- [DHCPv4 Individual Leasequery | 1118](#)
- [DHCPv6 Individual Leasequery | 1119](#)

Starting in Junos OS Release 16.1, subscriber management supports the individual leasequery feature, which enables the DHCPv4 or DHCPv6 relay agent to quickly and efficiently obtain the current lease information from a DHCP local server. The relay agent can lose locally stored lease information for various reasons, such as because the relay agent device was rebooted. When the relay agent subsequently receives data traffic from a client for forwarding, it no longer has the information to do so. A leasequery interaction with the local server can restore the information so that the relay agent can properly service its clients.

To configure individual leasequery operations, you enable support on both the DHCP relay agent and the DHCP server. You can configure details of the communication between the relay agent and the server. You must issue the `request dhcp leasequery` or `request dhcpv6 leasequery` command to trigger the relay agent to send the query.

By default, the relay agent sends the query to all known local servers. You can limit the servers it communicates with by specifying a server address or a named group of servers. You can also limit the query to servers in a particular logical system, routing instance, or LS:RI combination.

DHCPv4 Individual Leasequery

The DHCPv4 leasequery can be one of several types, a query by address, client ID, or MAC address. You determine the query type when you trigger the query by issuing the `request dhcp relay leasequery` command. You specify that the DHCPv4 relay agent includes in the DHCPLEASEQUERY message one of the following values to enable the local server to identify the binding information requested by the agent:

- IP address of a client lease—The local server returns binding information for the most recent client that was assigned that IP address.
- Client identifier of the client device—The local server returns binding information for the IP address that was most recently used by a client that has the specified client identifier (option 61). The identifier is unique across the server's administrative domain. If that client has accessed other IP addresses through this server, then the server returns a list of those addresses in the associated IP option (option 92).
- MAC address of the client device—The local server returns binding information for the most recent client that has that MAC address. If that client has accessed other IP addresses through this server, then the server returns a list of those addresses in the associated IP option (option 92).

The DHCP relay agent includes the parameter request list option (option 55) in the DHCPLEASEQUERY message. This list includes specific options related to the binding information for the IP address returned by the local server. For example, the request list typically includes the relay agent information option (option 82). The local server includes the requested information in a DHCPLEASEACTIVE sent to the relay agent.

The DHCPLEASEACTIVE message includes the client last transaction time option (option 91). The value of this option is the interval in seconds between when the IP address was most recently used in an interaction between the client and server and the time the server sends DHCPLEASEACTIVE message. For example, if the last interaction was at 08:00:00 and the message is sent at 09:00:00, then the option value is 3600.

[Table 36 on page 1119](#) describes the message types for DHCPv4 individual leasequery.

Table 36: DHCPv4 Individual Leasequery Message Types

Message Type	Option 53 Type Value	Description
DHCPLEASEQUERY	10	Sent by the relay agent to the DHCP local server to restore information.
DHCPLEASEUNASSIGNED	11	Response from the local server when the IP address associated with the client is controlled by the server but is not currently leased. This response is sent only for a query by IP address.
DHCPLEASEUNKNOWN	12	Response from the local server when the server has no knowledge of the information in the query.
DHCPLEASEACTIVE	13	Response from the local server when it has leased an address to the client. The response includes full binding information about that address.

DHCPv6 Individual Leasequery

The query type is conveyed in the LQ_Query option (option 44). The DHCPv6 relay agent query type can be by address or by client ID. You determine the query type when you trigger the query by issuing the request `dhcpv6 relay leasequery` command. You specify that the DHCPv6 relay agent includes in the LEASEQUERY message one of the following values in the option request option (option 6) to enable the local server to identify the binding information requested by the agent:

- IPv6 address of a client lease—The local server returns binding information for the most recent client that is bound to that address or has been delegated a prefix that contains the address. The query-options field in option 44 includes the IAADDR option (option 5).

- DHCP unique identifier (DUID) of the client device—The local server returns binding information for the IP address that was most recently used by a client that has the specified DUID. The DUID is the IPv6 identifier for the client. The identifier is unique across the server's administrative domain. The local server can return a list of addresses if the client is found on more than one link address. The query-options field in option 44 includes the Client Identifier option (option 1).

The query-options field in option 44 can also include the option request option (option 6) to list DHCPv6 option codes for specific information desired from the local server for each client.

The LEASEQUERY-REPLY message includes the client data option (option 45) to provide information for a single client on a single link. This information is conveyed as DHCPv6 options in the client-options field. Option 45 includes the following options as a minimum and any other options requested by the relay agent in the LEASEQUERY option request option (option 6):

- Client Identifier (option 1)—DUID that identifies the DHCPv6 client.
- IAADDR (option 5)—Address in an identity association for temporary addresses (IA_TA) or nontemporary addresses (IA_NA). Can be included with the IAPREFIX option.
- IAPREFIX (option 26)—Prefix in an identity association for prefix delegation (IA_PD). Can be included with the IAADDR option.
- CLT option (option 46)—The time in seconds since the server last interacted with the client on that link. This option corresponds to the DHCPv4 client last transaction time option.

The following options are examples of additional options that can be included in the LEASEQUERY-REPLY message:

- LQ relay data option (option 47)—The complete relay agent information that was used when the client last communicated with this server. The local server returns this option only when it is requested in the LEASEQUERY options request option (option 6).
- LQ client link option (option 48)—Identifies the link addresses on which the client has at least one binding. The LEASEQUERY-REPLY message includes this option when both of the following are true: the LEASEQUERY does not specify a link address and the client is found on more than one link. When the relay agent receives this information, it can submit a new LEASEQUERY for each address listed in option 48.

[Table 37 on page 1121](#) describes the message types for DHCPv6 individual leasequery.

Table 37: DHCPv6 Individual Leasequery Message Types

Message Type	DHCPv6 Type Value	Description
LEASEQUERY	14	Sent by the relay agent to the DHCP local server to restore information. Includes the LQ option (option 44) to specify the type of query, a link address, and any particular option information needed from the local server.
LEASEQUERY-REPLY	15	Response from the local server when the IP address associated with the client is controlled by the server but is not currently leased. This response is sent only for a query by IP address.

The LEASEQUERY-REPLY message sent by the DHCPv6 local server can return the status code option (option 13) to provide information about the status of the query. [Table 38 on page 1121](#) lists the status codes.

Table 38: DHCPv6 Individual Leasequery Status Codes

Code	Status	Description
7	UnknownQueryType	The server does not recognize or does not support the query.
8	MalformedQuery	The query is not valid; for example it might be missing a required option.
9	NotConfigured	The local server does not have the required address in its configuration.
10	NotAllowed	The local server does not allow the relay agent to send this query type.

DHCP Bulk Leasequery

IN THIS SECTION

- [DHCPv4 Bulk Leasequery | 1122](#)
- [DHCPv6 Bulk Leasequery | 1124](#)

Starting in Junos OS Release 16.1, subscriber management supports the bulk leasequery feature, which enables each request from the DHCP relay agent to retrieve lease information for multiple subscribers in bulk from a configured DHCP server in a programmed manner. Bulk leasequery is more resource-efficient than using multiple individual leasequeries to gather the same information. This is particularly useful in scaled environments with thousands of clients per relay agent.

Bulk leasequery uses a TCP connection between the DHCP relay agent and a configured DHCP server in the same logical system/routing instance. The TCP connection is more reliable and consumes fewer resources than the UDP connection used for the individual leasequery process. Bulk leasequery also extends the individual leasequery by providing additional query options and functionality.

To configure bulk leasequery operations, you enable support on both the DHCP relay agent and the DHCP server. You can configure details of the communication between the relay agent and the server. You must issue the `request dhcp bulk-leasequery` or `request dhcpv6 bulk-leasequery` command to trigger the relay agent to send the leasequery.

By default, the relay agent sends the query to all known local servers. You can limit the servers it communicates with by specifying an address for a server or a named group of servers. You can also limit the query to servers in a particular logical system, routing instance, or LS:RI combination.

DHCPv4 Bulk Leasequery

For DHCPv4 bulk leasequery, the DHCPv4 relay agent opens a TCP connection through port 67 to the DHCPv4 local server. When the connection is established, the relay agent sends a DHCPBULKLEASEQUERY message to the server. The query can contain any one of the following to enable the local server to identify the information needed by the agent:

- All configured IP addresses—The local server returns binding information for all IP addresses configured in the local server. The information is returned regardless of whether the IP addresses are part of a currently active binding. This enables the relay agent to update its database with all address changes that occurred after some point in time.
- Client identifier of the client device—The local server returns binding information for the IP address that was most recently used by a client that has the specified client identifier (option 61). The identifier is unique across the server's administrative domain.



NOTE: Unlike individual leasequery, the server does not use the associated IP option (option 92) to return a list of other IP addresses that the client has accessed through this server. Instead the server returns binding information for all these IP addresses

- MAC address of the client device—The local server returns binding information for the most recent client that has that MAC address.



NOTE: Unlike individual leasequery, the server does not use the associated IP option (option 92) to return a list of other IP addresses that the client has accessed through this server. Instead the server returns binding information for all these IP addresses

- Relay agent identifier—The local server returns binding information for all currently active leases assigned to the client that has the specified relay agent identifier (Option 82, suboption 12). The identifier is unique across the server's administrative domain.
- Remote ID of an access circuit used by the client to identify the circuit to the DHCP client—The local server returns binding information for all currently active leases assigned to clients that use that Agent Remote ID (option 82, suboption 2). This query is particularly useful in scaled environments with thousands of clients per relay agent. The other queries do not return consolidated lease information for all clients on a circuit.

The DHCPv4 local server replies to the relay agent with the same DHCPLEASEACTIVE and DHCPLEASEUNASSIGNED messages used for individual leasequery, as described in ["DHCPv4 Individual Leasequery Message Types" on page 1116](#). Each message corresponds to a single binding identified by the query.

When the server has returned all the bindings associated with the request, it sends a DHCPLEASEQUERYDONE message to the relay agent. If a connection is lost while processing a bulk leasequery, DHCP cannot determine how much of the requested information the relay agent received before the connection went down. Consequently, the relay agent must retry the query.

For any of the query methods, the DHCP relay agent can include the following qualifier:

- query-start-time—Returns bindings that changed on or after the time specified in the query.
- query-end-time—Returns bindings that changed on or before the time specified in the query.

These query times enable an agent to recover only binding information that it lost since it last committed all its information to stable storage.

[Table 39 on page 1123](#) describes the message types specific to DHCPv4 bulk leasequery.

Table 39: DHCPv4 Bulk Leasequery Message Types

Message Type	Option 53 Type Value	Description
DHCPBULKLEASEQUERY	14	Sent by the relay agent to the DHCP local server to restore information.

Table 39: DHCPv4 Bulk Leasequery Message Types (Continued)

Message Type	Option 53 Type Value	Description
DHCPLEASEQUERYDONE	15	Response from the local server when it has returned all binding information associated with the bulk request.

The messages sent by the DHCPv4 local server can return the status code option (option 151) to provide information about the status of the query. In DHCPLEASEACTIVE and DHCPLEASEUNASSIGNED messages, the code corresponds to the status for the individual binding request. In DHCPLEASEQUERYDONE messages, the code corresponds to the bulk leasequery request as a whole. [Table 40 on page 1124](#) lists the status codes.

Table 40: DHCPv4 Bulk Leasequery Status Codes

Code	Status	Description
0	Success	The request has been successfully completed. The absence of option 151 also indicates success.
1	UnSpecFail	The request failed for an unspecified reason.
2	QueryTerminated	The local server either could not perform the query or it terminated the query early. In the latter case, a text string indicates the cause.
3	MalformedQuery	The query was not understood by the local server.
4	NotAllowed	The query was understood but not allowed.

DHCPv6 Bulk Leasequery

For DHCPv6 bulk leasequery, the DHCPv6 relay agent opens a TCP connection through port 67 to the DHCPv6 local server. When the connection is established, the relay agent sends a LEASEQUERY message to the server. The query type is conveyed in the LQ_Query option (option 44). The query type can be any one of the following to enable the local server to identify the information needed by the agent:

- All configured IP addresses—The local server returns binding information for all IP addresses configured in the local server. The information is returned regardless of whether the IP addresses are part of a currently active binding. This enables the relay agent to update its database with all address changes that occurred after some point in time.
- Client identifier of the client device—The local server returns binding information for the IP address that was most recently used by a client that has the specified client identifier (option 61). The identifier is unique across the server's administrative domain.



NOTE: Unlike individual leasequery, the server does not use the associated IP option (option 92) to return a list of other IP addresses that the client has accessed through this server. Instead the server returns binding information for all these IP addresses

- MAC address of the client device—The local server returns binding information for the most recent client that has that MAC address.



NOTE: Unlike individual leasequery, the server does not use the associated IP option (option 92) to return a list of other IP addresses that the client has accessed through this server. Instead the server returns binding information for all these IP addresses

- Relay agent identifier—The local server returns binding information for all currently active leases assigned to the client that has the specified relay agent identifier (Option 82, suboption 12). The identifier is unique across the server's administrative domain.
- Remote ID of an access circuit used by the client to identify the circuit to the DHCP client—The local server returns binding information for all currently active leases assigned to clients that use that Agent Remote ID (option 82, suboption 2). This query is particularly useful in scaled environments with thousands of clients per relay agent. The other queries do not return consolidated lease information for all clients on a circuit.

For a DHCPv6 bulk leasequery, you can optionally specify the `trigger automatic` option to configure the DHCPv6 relay agent to automatically initiate the bulk leasequery operation whenever the `jdhcpd` process starts a connection with the session database (SDB) and no bound subscribers are present in the database. For example, the automatic process would ensure that the bulk leasequery always updates the DHCP relay information after a reboot, GRES, or ISSU operation, and if there are no bound subscribers.

DHCPv6 bulk leasequery uses the LEASEQUERY and LEASEQUERY-REPLY messages used by DHCPv6 individual leasequery, but their behavior and meaning is slightly different for bulk leasequery. [Table 41 on page 1126](#) lists these messages and describes two other message types are specific to DHCPv6 bulk leasequery.

Table 41: DHCPv6 Bulk Leasequery Message Types

Message Type	DHCPv6 Type Value	Description
LEASEQUERY	14	Sent by the relay agent to the DHCP local server to restore information.
LEASEQUERY-REPLY	15	<p>Response from the local server to indicate the success or failure of the query. It also conveys information, like the server Id and client ID, that does not change in the context of a single query and reply.</p> <p>When the query is successful, only a single LEASEQUERY-REPLY is returned. This message also includes the binding information for the first client. Additional binding data is returned in the LEASEQUERY-DATA message.</p> <p>When the query fails, a single LEASEQUERY-REPLY is returned with no binding information.</p>
LEASEQUERY-DONE	16	<p>Response from the local server that indicates the end of a group of related leasequery replies. A single LEASEQUERY-DONE message is sent after all replies to the request have been sent to the relay agent.</p> <p>The TCP connection between the relay agent and server is closed when this message is received.</p>
LEASEQUERY-DATA	17	<p>Response from the local server with information about the leases for a single DHCPv6 client or about prefix delegation bindings on a single link.</p> <p>This message is sent only when the bulk leasequery returns data for multiple clients. In this case, the LEASEQUERY-REPLY message conveys information for the first client, then a LEASEQUERY-DATA message is sent for each of the other clients.</p>

The messages sent by the DHCPv6 local server can return the status code option (option 13) to provide information about the status of the query. In LEASEQUERY-REPLY messages, the code corresponds to the status for the individual binding request. In LEASEQUERY-DONE messages, the code corresponds to the bulk leasequery request as a whole. LEASEQUERY-DATA messages do not include a status code. DHCPv6 bulk leasequery supports the DHCPv6 individual leasequery status codes listed in ["DHCPv6 Individual Leasequery Status Codes" on page 1116](#). The messages can also include the status code added for bulk leasequery described in [Table 42 on page 1127](#).

Table 42: DHCPv6 Bulk Leasequery Status Code

Code	Status	Description
11	QueryTerminated	The local server cannot perform a query or it has prematurely terminated the query for some reason. For example, the local server is being shut down or has insufficient resources to collect the requested information.

DHCP Active Leasequery

IN THIS SECTION

- [DHCPv4 Active Leasequery | 1128](#)
- [DHCPv6 Active Leasequery | 1130](#)
- [Chassis-Level Redundancy with Active Leasequery | 1131](#)
- [Interface-Level Redundancy with Active Leasequery Topology Discovery | 1133](#)

Starting in Junos OS Release 19.1R1, DHCP active leasequery addresses the situation where it is desirable for the relay agent to receive periodic updates of client information to keep up with dynamic DHCP binding activity. Individual and bulk leasequery provide information only when it is requested; if the client information is later updated on the local server, that information is not passed to the relay agent unless the relay agent sends another query to the local server.

Active leasequery enables servers to provide live updates of client information whenever the binding state changes. You can optionally configure active leasequery to send the live updates of binding information to multiple relay agent peers, supporting relay agent chassis-level redundancy. Live updating is initiated when the relay agent starts a TCP connection with a server or relay agent peer and sends the ACTIVELEASEQUERY message to indicate that the connection must stay open.

DHCP does not close the TCP connection unless certain conditions occur, mostly related to the configurable timeout or idle timeout periods:

- When a connection request is received in a logical system or routing instance that is not configured for active leasequery.
- When the connection is blocked during TCP read/write operations long enough for the timeout period to expire, the connection is closed and can be restarted. The read operation is when the relay agent is trying to read replies to the query. The write operation is when the server or peer relay agent is trying to send replies to a relay agent

- When no traffic is received on the connection for the duration of the idle timeout period.

During active leasequery operations, binding information is updated only when it changes. Consequently, there are periods during which the server or peer relay agent sends no information. If the period is longer than the idle-timeout, the connection is dropped. To avoid inappropriate connection drops, the server or peer relay agent sends DHCPLEASEACTIVE (DHCPv4) or LEASEQUERY-DATA (DHCPv6) messages at intervals equal to one-half of the idle timeout period. These messages contain no binding information because they are sent when no updates are available. These messages keep the connection alive by serving as hello or keepalive messages signaling that the lack of activity is not a problem.

When the TCP connection does close, the relay agent tries to reestablish the connection. The retry attempts include an option that signals the server or peer relay agent to send binding information that changed from the time that the TCP connection shut down. This information is sometimes referred to as the catch-up information. The option specifies the absolute timestamp when the connection shut down; that is, the time of the last successful communication with the server or peer relay agent. DHCPv4 uses the query-start-time option (option 154). DHCPv6 uses the LQ_START_TIME option (option 101).

In some cases, the server or peer relay agent does not have all the information for binding changes since the timestamp. For example, the device might not have sufficient memory to store it all. In these cases, the device returns a DHCPLEASEQUERYSTATUS (DHCPv4) or LEASEQUERY-REPLY (DHCPv6) message is sent with a status code of DataMissing (5).



NOTE: Before you configure active leasequery, you must first configure bulk leasequery, because active leasequery uses the bulk leasequery mechanism. The active leasequery configuration fails commit check if bulk leasequery is not configured.

To configure active leasequery operations, you enable support on both the DHCP relay agent and the DHCP server. You can configure details of the communication for both the relay agent and the local server. Unlike individual and bulk leasequery, active leasequery has no query types. You do not trigger active leasequery with a request command. Instead, the trigger is automatic when active leasequery is configured.

DHCPv4 Active Leasequery

For DHCPv4 active leasequery, the DHCPv4 relay agent opens a TCP connection through port 67 to the DHCPv4 local server. When the connection is established, the relay agent sends a DHCPACTIVELEASEQUERY message to the server. The message signals that this is a long-term connection. It closes only as a result of a timeout.

The DHCPv4 local server replies to the relay agent with the same DHCPLEASEACTIVE and DHCPLEASEUNASSIGNED messages used for individual leasequery, as described in "[DHCPv4 Individual Leasequery Message Types](#)" on page 1116. Each message corresponds to a single binding

identified by the query. The DHCP local server continues to send the response messages whenever the binding information changes. [Table 43 on page 1129](#) describes the message types specific to DHCPv4 active leasequery.

Table 43: DHCPv4 Active Leasequery Message Types

Message Type	Option 53 Type Value	Description
DHCPACTIVELEASEQUERY	16	<p>Sent by the relay agent to the DHCP local server to enable live updating of binding information on the relay agent whenever that information changes on the local server.</p> <p>Can also be sent between peer relay agents to provide hot standby redundancy for binding information.</p>
DHCPLEASEQUERYSTATUS	17	<p>Response from the local server when it has returned binding information associated with the request.</p> <p>Because the TCP connection is long-lived, this message is also sent regularly when the connections is idle (no binding updates being sent). In this case the message includes a ConnectionActive status code (6) to notify the relay agent that the connection is still up.</p>

The messages sent by the local server can return the status code option (option 151). In DHCPLEASEACTIVE and DHCPLEASEUNASSIGNED messages, the code corresponds to the status of the individual response. In DHCPLEASEQUERYSTATUS messages, the code corresponds to the message stream for the active leasequery request as a whole. DHCPv4 active leasequery supports the bulk leasequery status codes listed in ["DHCPv4 Bulk Leasequery Status Codes" on page 1116](#). The messages can also include the status codes added for active leasequery described in [Table 44 on page 1129](#).

Table 44: DHCPv4 Active Leasequery Status Codes

Code	Status	Description
5	DataMissing	The requested binding information is not available. For example, when the local server or peer does not have enough data as requested with the query-start-time option, this status code is sent immediately in a LEASEQUERY-REPLY message.
6	ConnectionActive	The TCP connection is still active.

Table 44: DHCPv4 Active Leasequery Status Codes (Continued)

Code	Status	Description
7	CatchUpComplete	The local server has sent all the saved data requested by the relay agent.

DHCPv6 Active Leasequery

For DHCPv6 active leasequery, the DHCPv6 relay agent opens a TCP connection through port 67 to the DHCPv4 local server. When the connection is established, the relay agent sends an ACTIVELEASEQUERY message to the server. The message signals that this is a long-term connection. It closes only as a result of a timeout.

The DHCPv6 local server replies to the relay agent with the same LEASEQUERY-REPLY, LEASEQUERY-DATA, and LEASEQUERY-DONE messages used for bulk leasequery. Each message corresponds to a single binding identified by the query. The DHCP local server continues to send the response messages whenever the binding information changes. [Table 45 on page 1130](#) lists these messages and the query message type that is specific to DHCPv6 active leasequery.

Table 45: DHCPv6 Active Leasequery Message Types

Message Type	DHCPv6 Type Value	Description
ACTIVELEASEQUERY	22	<p>Sent by the relay agent to the DHCP local server to enable live updating of binding information on the relay agent whenever that information changes on the local server.</p> <p>Can also be sent between peer relay agents to provide hot standby redundancy for binding information.</p>
LEASEQUERY-REPLY	15	<p>Response from the local server to indicate the success or failure of the query. It also conveys information, like the server Id and client ID, that does not change in the context of a single query and reply.</p> <p>When the query is successful, only a single LEASEQUERY-REPLY is returned. This message also includes the binding information for the first client. Additional binding data is returned in the LEASEQUERY-DATA message.</p> <p>When the query fails, a single LEASEQUERY-REPLY is returned with no binding information.</p>

Table 45: DHCPv6 Active Leasequery Message Types (Continued)

Message Type	DHCPv6 Type Value	Description
LEASEQUERY-DONE	16	<p>Response from the local server that indicates that the connection should be terminated.</p> <p>For example, the sever can send this with a QueryTerminated status code (11) when the server is being shut down.</p>
LEASEQUERY-DATA	17	<p>Response from the local server with information about the leases for a single DHCPv6 client or about prefix delegation bindings on a single link.</p> <p>This message is sent only when the leasequery returns data for multiple clients. In this case, the LEASEQUERY-REPLY message conveys information for the first client, then a LEASEQUERY-DATA message is sent for each of the other clients.</p>

The messages sent by the DHCPv6 local server can return the status code option (option 13). DHCPv6 active leasequery supports the individual leasequery and bulk leasequery status codes listed in ["DHCPv6 Individual Leasequery Status Codes" on page 1116](#) and ["DHCPv6 Bulk Leasequery Status Code" on page 1116](#), respectively. The messages can also include the status codes added for active leasequery described in [Table 46 on page 1131](#).

Table 46: DHCPv6 Active Leasequery Status Codes

Code	Status	Description
12	DataMissing	The requested binding information is not available.
13	CatchUpComplete	The local server has sent all the saved data requested by the relay agent.
14	NotSupported	The local server has sent all the saved data requested by the relay agent.


Chassis-Level Redundancy with Active Leasequery

You can use active leasequery to enable binding information to be synchronized between multiple DHCP relay agent peers. For simplicity, this discussion explains the behavior with only two peers. When a peer relay agent restarts or its device reboots, the other relay can take over and provide services to all the DHCP clients without a visible outage. When the peer relay agent comes up again, it reestablishes

the TCP connection with the active peer. The peers then synchronize binding information. [Figure 5 on page 1132](#) shows a simple DHCP topology to support relay agent redundancy with the following characteristics:

- Each DHCP client connects to both relay agents.
- Both relay agents connect to the same DHCP server.
- When you configure the active leasequery statement on each relay agent, you also specify the other relay agent as a peer.
- The peers use the same active leasequery messages for communication as explained in [Table 43 on page 1129](#) and [Table 45 on page 1130](#). Although it is not shown here, when an external RADIUS server is part of the topology, there are no differences in interactions with the RADIUS server.

Figure 5: Simple Topology for DHCP Redundancy with Active Leasequery



The following sequence describes how the relay agents establish the peer relationship and share binding information when active leasequery is configured on both. This example is for DHCPv4, but the mechanism is the same for DHCPv6.

1. Both relay agents have active DHCP client bindings, but active leasequery is not yet configured.
2. You configure active leasequery on both relay agents, specify each other as peers, and commit the configuration.
3. Both peer agents attempt to establish a TCP connection when the configuration is committed. Suppose relay agent Relay Agent 1 successfully establishes the connection. The attempt from peer Relay Agent 2 is dropped.
4. Relay Agent 1 then sends an ACTIVELEASEQUERY message to Relay Agent 2.
5. Relay Agent 2 sends information about the bindings in its subscriber database to Relay Agent 1. It also sends its own ACTIVELEASEQUERY message to Relay Agent 1 to collect the peer's client information.

6. Relay Agent 1 sends its binding information to Relay Agent 2. Relay Agent 1 and Relay Agent 2 each process the received binding information and commit it to their respective databases.
7. As each relay agent updates binding information for its own clients—such as license renewals, new requests, lease expirations and so on—it sends a leasequery response message with the updated information to its peer when each change occurs.
8. Now suppose Relay Agent 1 is rebooted. The TCP connection drops. Relay Agent 2 tries to reestablish the connection with Relay Agent 1. In the meantime, the DHCP subscriber traffic that used to flow through Relay Agent 1 now flows through Relay Agent 2 without interruption.
9. Active leasequery is triggered on Relay Agent 1 when it comes back up. The TCP connection is reestablished and the peers exchange ACTIVELEASEQUERY messages. Relay Agent 1 has no binding information to share at this point. Relay Agent 2 sends all of its current binding information to Relay Agent 1; this information might have changed while Relay Agent 1 was out of service. The result is that both relay agents now have synchronized databases.

Interface-Level Redundancy with Active Leasequery Topology Discovery

Starting in Junos OS Release 19.2R1, topology discovery enables DHCP relay peers to discover information about each other's subscriber interfaces. Topology discovery is necessary in a network topology with an M:N subscriber group redundancy configuration. In this configuration, a BNG that hosts a DHCP relay agent acts as the primary router for a subscriber redundancy group. The primary router handles traffic for the subscriber redundancy group. One or more other BNGs that host peer relay agents serve as backups for subscriber redundancy groups on the primary.

A particular BNG can be the backup for multiple subscriber redundancy groups, but each redundancy group is backed up to only one BNG. If the primary BNG fails, the backup BNG for each subscriber redundancy group that is affected by the failure is elected as the new primary for that redundancy group. The new primary continues to serve the subscriber redundancy group seamlessly and without disruption. See "[M:N Subscriber Redundancy Overview](#)" on [page 1511](#) for more information about M:N redundancy.

Interface-level subscriber redundancy is based on the logical interface for the access link. In this situation, the interface name of the access interface for a subscriber redundancy group does not need to be the same on the primary and backup peers. This behavior is different than that for chassis-level relay agent redundancy, where the access interface names must be identical on the relay agent peers.

Because the interface names can be different for the primary and backup relay agents, DHCP needs to discover the relationship between the interface for each subscriber redundancy group on the primary and the corresponding interface on the backups. Topology discovery provides that information.

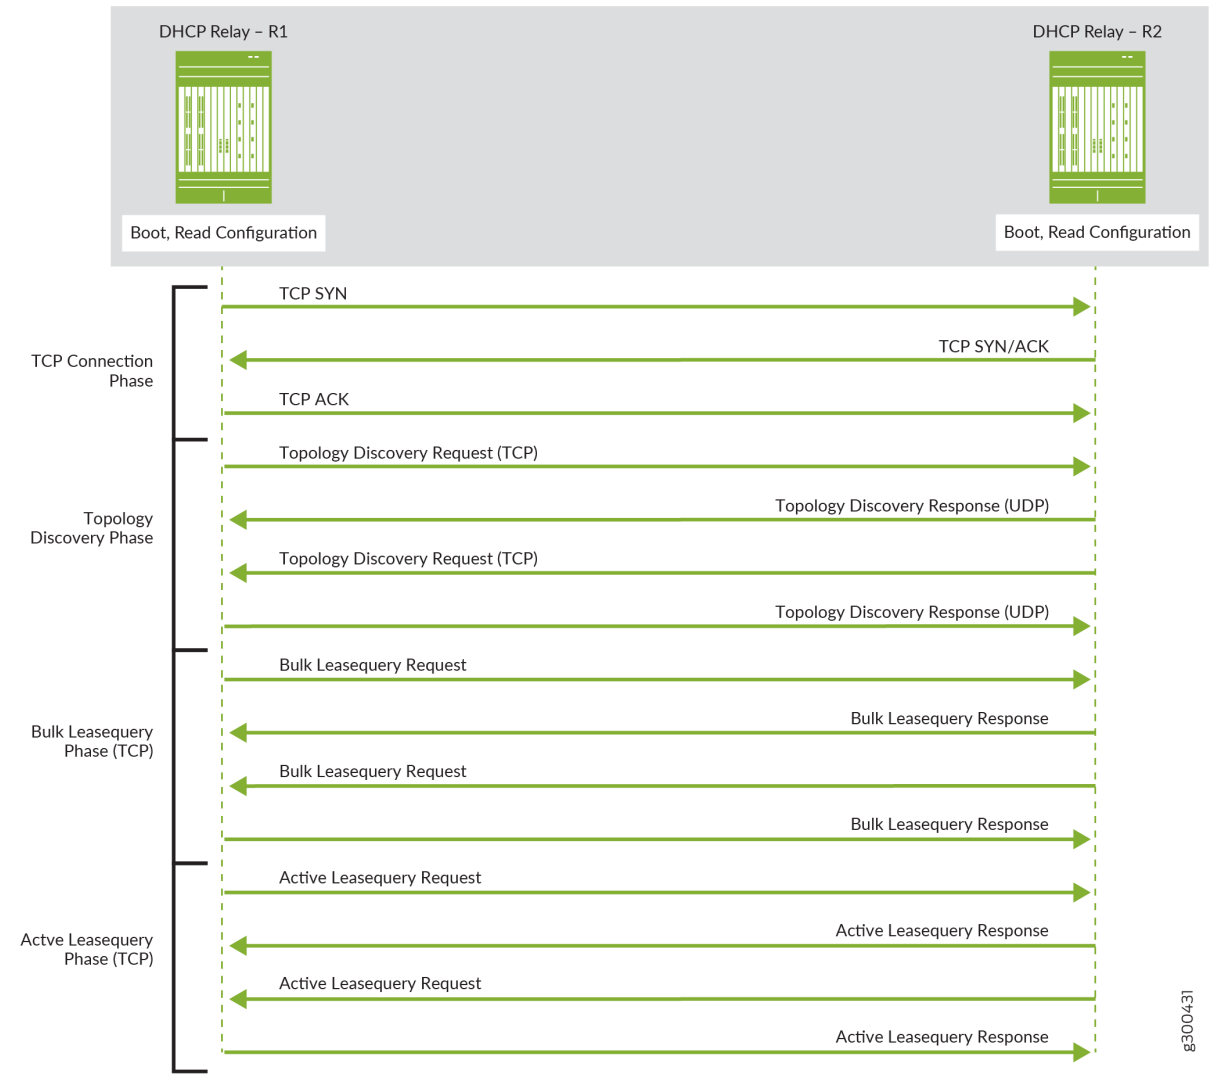
Topology discovery enables the primary and backup relay agents to automatically build a translation table that maps the local and remote access interfaces for each subscriber redundancy group. If the primary fails, then the backup elected to be the new primary uses its translation table to immediately

manage the subscriber redundancy groups affected by the failure. The failover itself is transparent to the DHCP clients associated with the subscriber redundancy groups.

Topology discovery is an active leasequery option. Active leasequery enables the peers to synchronize the binding information for subscribers in the subscriber redundancy groups corresponding to the interfaces added to the translation table. DHCP translates the binding information to use the local interface on the backup instead of the interface on the primary.

When you configure topology discovery, the entire DHCP leasequery process consists of four connection phases, as shown in [Figure 6 on page 1134](#).

Figure 6: Topology Discovery Connection Phases



1. TCP connection phase—A TCP connection is established between the peer relay agents.

2. **Topology discovery phase**—The peers exchange topology discovery messages to determine the matching access interfaces for each subscriber redundancy group on the peers. The remote peer matches an interface based on the VLAN ID and subnet. Each peer sends a query for all of its access interfaces and receives a response, so that all peers can build a translation table of connected local and remote interface pairs for the subscriber redundancy groups.
3. **Bulk leasequery phase**—The peers establish the bulk leasequery relationship required for active leasequery to operate. Bulk leasequery enables the relay agents to retrieve lease information for multiple subscribers from a configured DHCP server in bulk rather than in a series of individual queries and responses. In this phase DHCP collects in bulk all the binding information for the first time.
4. **Active leasequery phase**—Active leasequery ensures that binding information is synchronized whenever it changes, without the need for subsequent queries. The primary relay agent sends the bindings relative to its local Agent Circuit ID (the name of the access interface). The backup relay agent uses its translation table to obtain the corresponding Agent Circuit ID on the backup to install the subscribers.

To restrict the information that is synchronized to only the subscribers that use a particular access interface—in other words, a subscriber redundancy group, active leasequery uses the query by giaddr (DHCPv4) or linkaddr (DHCPv6) method when you configure topology discovery. The gateway IP address (giaddr or linkaddr) is what a relay agent uses to determine where to send information downstream. The value of the giaddr is the access interface. The relay agent evaluates the giaddr/linkaddr and sends information to the DHCP client that uses the access interface matching the giaddr/linkaddr.

What this means for subscriber redundancy is that by using the giaddr/linkaddr query, active leasequery requests only information for subscribers on that access interface. Consequently, it synchronizes only that subscriber information from the primary relay agent to the backup relay agent. This is a much smaller set of subscribers than if the active leasequery used the query by relay-id method, which returns information for all subscribers on the entire chassis.

The result of this process is that each peer agent installs the subscribers for each redundancy group it handles. When the primary relay agent fails over, the backup already has the necessary subscriber information to maintain the affected subscriber sessions without interruption.



NOTE: The bulk leasequery and active leasequery connection phases run over the TCP connection. In contrast, during the topology discovery phase, DHCP sends the query messages over TCP, but sends the topology discovery response messages over UDP. The TCP path can be anything, but the UDP path must be through the access interface; this is how the peers confirm their access interfaces are connected.

Topology Discovery Messages

Topology discovery uses the standard individual leasequery messages. For DHCPv4, these are DHCPLEASEQUERY and DHCPLEASEACTIVE. For DHCPv6, these are LEASEQUERY and LEASEQUERY-REPLY. The difference that makes these messages specifically topology discovery messages is that each message includes a proprietary suboption value in the vendor-specific option (option 43 for DHCPv4 and option 17 for DHCPv6). The proprietary value is a string, topology_discover_lq. [Table 47 on page 1136](#) lists the information carried in the query and reply messages.



NOTE: Topology discovery for VRRP M:N redundancy uses TCP for the query and UDP for the response. Topology discovery for pseudowire M:N redundancy uses TCP for both query and response.

Table 47: Information Carried in Topology Discovery Query and Response Messages

Query	Response
Transaction ID (xid)—This number is unique per chassis. DHCP generates the xid for an access interface used by a subscriber redundancy group. The xid is carried in the DHCP header.	Transaction ID (xid)—The same value received in the request message.
Client identifier (DHCPv4 option 61; DHCPv6 option 1)—A string that identifies the DHCP client, based on the LACP MAC address.	Client identifier (DHCPv4 option 61; DHCPv6 option 1) —The same value received in the request message.
n/a	Server identifier (DHCPv4 option 54; DHCPv6 option 2)—A string that identifies the relay agent, based on the LACP MAC address
Agent Circuit ID (DHCPv4 option 82; DHCPv6 option 18)—Interface name of the access interface for which the query is made. This is used for translating local and peer interface ID.	Agent Circuit ID (DHCPv4 option 82; DHCPv6 option 18)—Interface name of the matching access interface on the peer. This is used for translating local and peer interface ID.

Table 47: Information Carried in Topology Discovery Query and Response Messages (*Continued*)

Query	Response
<p>Vendor Specific Option (DHCPv4 option 43; DHCPv6 option 17)—This option carries the following information specific for the vendor, Juniper Networks:</p> <ul style="list-style-type: none"> • Suboption 1—A string with the value <code>topology_discover_lq</code>. This is proprietary and makes the message a topology discovery message. • Suboption 2—IP (subnet) address of the querying interface. This is the address that the DHCP relay agent puts in the <code>giaddr</code> field in messages it sends to the DHCP server. • Suboption 3—Subnet mask of the querying interface. • Suboption 4—VLAN ID of the querying interface. • Suboption 5—Logical system/routing instance of the querying interface in the format <i>logical-system-name;routing-instance-name</i>. • Suboption 6—Shared common key of the querying interface. This is an ASCII string of up to 63 characters. 	<p>Vendor Specific Option (DHCPv4 option 43; DHCPv6 option 17)—This option carries the following information:</p> <ul style="list-style-type: none"> • Suboption 1—A string with the value <code>topology_discover_lq</code>. This is proprietary and makes the message a topology discovery message. • Suboption 2—IP address of the matching interface on the peer. • Suboption 3—Subnet mask of the matching interface on the peer. • Suboption 4—VLAN ID of the matching interface on the peer. • Suboption 5—Logical system/routing instance of the matching interface on the peer. • Suboption 6—Shared common key of the matching interface on the peer. The same value received in the request message. <p>For M:N redundancy using VRRP, matching is based on the querying interface's name and subnet address, VLAN ID, and transaction ID received in the request.</p> <p>For M:N redundancy using pseudowires, matching is based on the querying interface's shared common key and transaction ID received in the request.</p>

The peer relay agents exchange topology discovery messages when any of the following occurs:

- You configure a new peer relay agent.
- The router restores an access interface connection so that the link is up.
- The router starts up.
- The `jdhcpd` process restarts.
- You configure active leasequery.

- The topology changes. The relay agent detects this change when a topology discovery query arrives on a link that was previously discovered.

For a detailed explanation of how topology works with M:N subscriber redundancy, see "[M:N Subscriber Redundancy Overview](#)" on page 1511.

Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations

When configuring individual, bulk, or active leasequery support, consider the following guidelines:

- The router supports simultaneous configuration of individual leasequery, bulk leasequery, and active leasequery. Active leasequery requires bulk leasequery to be configured.
- The router supports simultaneous dual-stack configuration for both DHCPv4 and DHCPv6. However, for dual stack environments, you must trigger the DHCPv4 and DHCPv6 individual leasequery or bulk leasequery operations separately.
- DHCP relay agent supports individual leasequery or bulk leasequery over static and dynamic interfaces. Active leasequery is supported only on server-facing static interfaces or peer-facing static interfaces for chassis redundancy.
- DHCP local server supports bulk leasequery only on relay-facing static interfaces.
- DHCP local server listens for bulk leasequery and active leasequery requests from the DHCP relay agent on the TCP connection on port 67 for DHCPv4 and on port 547 for DHCPv6.
- Bulk leasequery and active leasequery are not supported for DHCP over PPP/PPPoE.
- Active leasequery is supported over the following stack combinations:
 - DHCP over static interfaces (ge/ae/xr/irb/ps) (Support for ps interfaces added in Junos OS Release 20.1R1.)
 - DHCP over IP Demux interfaces
 - DHCP over VLAN Demux interfaces
 - DHCP over IP over VLAN Demux interfaces
- Starting in Junos OS Release 19.1R1, the DHCPv4 relay agent inserts the Relay-ID option in each packet it forwards to the DHCP local server as follows:
 - The relay agent always inserts the option in non-snooped packets.
 - The relay agent inserts the option in snooped packets only when bulk leasequery is configured in that LS:RI.

- If the network includes integrated routing and bridging (IRB) interfaces, you must configure DHCP relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82. DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

Configuring and Using DHCP Individual Leasequery

The individual leasequery operation updates a DHCP relay agent's lease database with information related to a single, specified subscriber. You identify DHCPv4 subscribers by the DHCP client's IPv4 address, MAC address, or client ID. You identify DHCPv6 subscribers by the DHCP client's IPv6 address or client ID.

Before you begin, read ["Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations" on page 1138](#) and ensure that the following required support is configured on the DHCP relay agent.

- (DHCPv4 only) DHCP relay agent inserts option 82 suboption 1 (Agent Circuit ID), in the DHCP packets that the relay forwards to DHCP servers. See ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).

If the network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCP relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82. DHCP relay agent uses the layer 2 interface name when restoring the lease database using leasequery or bulk leasequery.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option-82 circuit-id include-irb-and-l2
```

- (DHCPv4 only) DHCP relay agent always includes the new option 82 information in the DHCP packets that the relay forwards to DHCP servers. See ["Overriding Option 82 Information" on page 1078](#).

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

- (DHCPv6 only) DHCP relay agent inserts the DHCPv6 Interface-ID (option 18) in the packets that the relay forwards to DHCPv6 servers. See ["Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets" on page 1458](#).

If your network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCPv6 relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82.

DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id include-irb-and-l2
```

Use the following steps to configure and use the individual leasequery operation.

1. Configure DHCP relay agent to support leasequery:

Configure the leasequery parameters the DHCP relay agent uses when querying the DHCP local servers. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level.

- a. Specify that you want to configure leasequery options for the DHCP relay agent.

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

- b. Specify the number of seconds DHCP relay waits before resending leasequery messages to the configured DHCP servers in the same logical system/routing instance.

```
[edit forwarding-options dhcp-relay leasequery]
user@host# set timeout seconds
```

- c. Specify the number of times DHCP relay resends leasequery messages . DHCP relay resends the messages when the configured timeout value expires. The messages are resent if the DHCP relay has not received confirmed lease information for a client.

```
[edit forwarding-options dhcp-relay leasequery]
user@host# set attempts number-of-attempts
```

2. Configure DHCP local server to support leasequery:

Configure the leasequery parameters the DHCP local server uses when responding to leasequery messages from a DHCP relay agent. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit system services dhcp-local-server dhcpv6] hierarchy level.

- a. Enable leasequery support for the DHCP local server.

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

- b. (Optional) Specify that the DHCP local server responds to a leasequery by sending the binding information only to restricted requestors. For DHCPv4, restricted requestors are those whose giaddr matches the giaddr of the client. For DHCPv6, the client ID of the request must match the relay ID of the client. This step provides additional security by ensuring that the requestor is the originator of the binding request.

```
[edit system services dhcp-local-server allow-leasequery]
user@host# set restricted-requestor
```

3. Initiate the leasequery operation on the DHCP relay agent. See ["Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 1152](#).

Use the supported `show` and `clear` commands to manage and display information about the bulk leasequery operation for the DHCP relay agent and the DHCP local server. See ["Verifying and Managing DHCP Individual and Bulk Leasequery Configurations" on page 1157](#).

Configuring and Using DHCP Bulk Leasequery

The bulk leasequery operation updates a DHCP relay agent's lease database with information for multiple subscribers, as opposed to the individual leasequery, which queries individual bindings for known targets only. Bulk leasequery also extends the individual leasequery by providing additional query options and functionality.

Before you begin, read ["Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations" on page 1138](#) and ensure that the following required support is configured on the DHCP relay agent.

- (DHCPv4 only) DHCP relay agent inserts option 82 suboption 1 (Agent Circuit ID), in the DHCP packets that the relay forwards to DHCP servers. See ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).

If the network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCPv6 relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82.

DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option-82 circuit-id include-irb-and-l2
```

- (DHCPv4 only) DHCP relay agent always includes the new option 82 information in the DHCP packets that the relay forwards to DHCP servers. See ["Overriding Option 82 Information" on page 1078](#).

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

- (DHCPv6 only) DHCP relay agent inserts the DHCPv6 Interface-ID (option 18) in packets forwarded to DHCPv6 servers. See ["Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets" on page 1458](#).

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id
```

If your network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCPv6 relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82. DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id include-irb-and-l2
```

Use the following steps to configure and use the bulk leasequery operation.

1. (Optional) Configure the number of connections that the router can use for bulk leasequery. Specify the maximum number of TCP connections the DHCP local server can simultaneously accept for bulk leasequery operations, and the number of simultaneous connections that the DHCP relay

agent can request for bulk leasequery. This is a chassis-wide configuration and includes all logical systems/routing instances, and all address families.

```
[edit system processes dhcp-service]
user@host# set accept-max-tcp-connections max-tcp-connections
user@host# set request-max-tcp-connections max-tcp-connections
```

2. Configure DHCP relay agent to support bulk leasequery:

Configure the bulk leasequery parameters the DHCP relay agent uses when querying the DHCP local servers. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level.

- a. Specify that you want to configure bulk leasequery options for the DHCP relay agent.

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

- b. Specify the number of seconds DHCP relay waits before retrying the TCP connection to send bulk leasequery messages to the configured DHCP servers in the same logical system/routing instance.

```
[edit forwarding-options dhcp-relay bulk-leasequery]
user@host# set timeout seconds
```

- c. Specify the number of times DHCP relay attempts the TCP connection with the local server to send bulk leasequery messages. DHCP relay resends the messages when the configured timeout value expires. The TCP connection is reestablished only to DHCP servers to which the connection failed or was abruptly closed.

```
[edit forwarding-options dhcp-relay bulk-leasequery]
user@host# set attempts number-of-attempts
```

- d. (Optional, DHCPv6 only) Specify the optional automatic trigger. The automatic trigger configures DHCPv6 relay agent to automatically initiate bulk leasequery whenever the jdhcpd process starts (for example, after a jdhcpd restart, a relay agent device reboot, a *graceful Routing Engine switchover*, or a unified ISSU) and there are no bound subscribers in the session database. The automatic bulk leasequery is always based on the relay agent Relay-ID option (option 53).



NOTE: When the automatic trigger support is configured, you can still use the CLI command to manually trigger bulk leasequeries separate from the automatic queries.

```
[edit forwarding-options dhcp-relay dhcpv6 bulk-leasequery]
user@host# set trigger automatic
```

3. Configure DHCP local server to support bulk leasequery:

Configure the parameters the DHCP local server uses when responding to bulk leasequery messages from a DHCP relay. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

- a. Enable bulk leasequery support for the DHCP local server.

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

- b. (Optional) Specify the maximum number of concurrent TCP connections allowed in the DHCP local server's logical system/routing instance:

```
[edit system services dhcp-local-server allow-bulk-leasequery]
user@host# set max-connections number-of-connections
```

- c. (Optional) Specify the maximum number of empty replies that the DHCP local server sends to a specific requestor. When the maximum number of replies is reached, the DHCP server closes the TCP connection to the requestor.

An empty reply is a response that contains no bindings or has an option status code error. Empty replies are often a response to an unauthorized requestor that has sent an invalid or incorrect query resulting in no binding. By limiting the number of empty replies that the DHCP local server sends, you prevent the connection from being taken over by unauthorized or malicious requestors.

```
[edit system services dhcp-local-server allow-bulk-leasequery]
user@host# set max-empty-replies number-of-replies
```

- d. (Optional) Specify that the DHCP local server sends the binding information to restricted requestors only. This step ensures that the requestor is the originator of the binding request.

For DHCPv4 leasequery and bulk leasequery requests, the giaddr of the requestor must match the giaddr of the client. For DHCPv6 bulk leasequery requests, the requestor's client ID in the bulk leasequery message must match the relay ID that was sent during binding creation.

```
[edit system services dhcp-local-server allow-bulk-leasequery]
user@host# set restricted-requestor
```

- e. (Optional) Specify the number of seconds that a connection on the TCP socket is idle before the DHCP local server closes the connection.

```
[edit system services dhcp-local-server allow-bulk-leasequery]
user@host# set timeout seconds
```

4. Initiate the bulk leasequery operation on the DHCP relay agent. See ["Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 1152](#).
 - Manually initiating bulk leasequery—(DHCPv6 only) Use the appropriate CLI command to manually initiate bulk leasequery. See ["Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 1152](#).
 - Automatically initiating bulk leasequery—When the automatic trigger feature is configured, DHCP relay agent initiates the bulk leasequery whenever the jdhcpd process starts and there are no bound subscribers in the session database.

Use the supported `show` and `clear` commands to manage and display information about the bulk leasequery operation for the DHCP relay agent and the DHCP local server. See ["Verifying and Managing DHCP Individual and Bulk Leasequery Configurations" on page 1157](#).

Configuring and Using DHCP Active Leasequery

Before you begin, read ["Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations" on page 1138](#) and ensure that the following required support is configured on the DHCP relay agent.

- (DHCPv4 only) DHCP relay agent inserts option 82 suboption 1 (Agent Circuit ID), in the DHCP packets that the relay forwards to DHCP servers. See ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).

If the network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCPv6 relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82.

DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option-82 circuit-id include-irb-and-l2
```

- (DHCPv4 only) DHCP relay agent always includes the new option 82 information in the DHCP packets that the relay forwards to DHCP servers. See ["Overriding Option 82 Information" on page 1078](#).

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

- (DHCPv6 only) DHCP relay agent inserts the DHCPv6 Interface-ID (option 18) in packets forwarded to DHCPv6 servers. See ["Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets" on page 1458](#).

If your network includes integrated routing and bridging (IRB) interfaces, you must also include the `include-irb-and-l2` statement, as shown in the following example. This statement configures DHCPv6 relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82. DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id include-irb-and-l2
```

- For chassis-level DHCP relay agent redundancy, the following guidelines apply:
 - The DHCP relay agent redundancy peers must all have identical subscriber configurations in order to have synchronized databases.
 - The complete interface names for the access interfaces (ge, xe, or ae) on which the subscribers come up must be identical on the DHCP relay agent redundancy peers.
- For interface-level DHCP relay agent primary/backup redundancy, the interface names do not have to be identical on the redundancy peers. The primary and backup relay agents use topology discovery to build translation tables that map local and remote (peer) interfaces for subscriber redundancy groups.



NOTE: When you configure topology discovery on all available logical interfaces, chassis-level redundancy is supported if the interface names and subscriber configurations match on the redundancy peers.

- Because active leasequery is an extension of bulk leasequery, you must configure bulk leasequery for active leasequery to operate. See ["Configuring and Using DHCP Bulk Leasequery" on page 1141](#).

The active leasequery operation sends live updates to DHCP relay agents for multiple subscribers when the DHCP binding information changes on the local server. You can also use active leasequery as part of a configuration to provide redundancy of binding information among peer relay agents.

Use the following steps to configure and use the active leasequery operation.



NOTE: These steps do not duplicate any of the bulk leasequery configuration. For example, the steps do not include configuring the maximum number of TCP connections, because that is part of the required bulk leasequery configuration.

1. Configure DHCP relay agent to support active leasequery:

Configure the active leasequery parameters the DHCP relay agent uses when querying the DHCP local servers.



NOTE: The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level.

- a. Specify that you want to configure active leasequery options for the DHCP relay agent.

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

- b. Specify the number of seconds DHCP relay waits when TCP read and write operations are blocked before terminating the TCP connection with the local server and then restarting it.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set timeout seconds
```

- c. Specify the number of seconds DHCP relay waits when no incoming data is received on the TCP connection before terminating the TCP connection with the local server and then restarting it.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set idle-timeout seconds
```

- d. (Optional) Specify the IP address for a peer with which this relay agent synchronizes information. The peer must also be configured for active leasequery.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set peer-address ip-address
```

- e. (Optional) Configure the relay agent to send topology discovery messages to determine the remote access interfaces for subscriber redundancy groups on similarly configured peer relay agents. Discovering the topology enables the relay agents to build translation tables of local and remote interfaces to support an interface-level, primary/backup redundancy scheme. See "[M:N Subscriber Redundancy Overview](#)" on page 1511 for information about using this type of redundancy.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set topology-discover
```

2. Configure DHCP local server to support active leasequery:

Configure the parameters the DHCP local server uses when responding to bulk leasequery messages from a DHCP relay. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit system services dhcp-local-server dhcpv6] hierarchy level.

- a. Enable bulk leasequery support for the DHCP local server.

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

- b. Specify the number of seconds DHCP local server waits when TCP read and write operations are blocked before terminating the TCP connection.

```
[edit system services dhcp-local-server allow-active-leasequery]
user@host# set timeout seconds
```

- c. (Optional) Specify the number of seconds that a connection on the TCP socket is idle before the DHCP local server closes the connection.

```
[edit system services dhcp-local-server allow-active-leasequery]
user@host# set idle-timeout seconds
```

3. Initiate the bulk leasequery operation on the DHCP relay agent. See ["Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 1152](#).



NOTE: There is no manual initiation for active leasequery. Active leasequery is automatic when both the following have occurred:

- Bulk leasequery has been configured and initiated.
- Active leasequery has been configured and committed.

Thereafter, DHCP relay agent automatically initiates active leasequery whenever the `jdhcpd` process starts (for example, after a reboot, a *graceful Routing Engine switchover*, or a unified ISSU) and when no bound subscribers are present in the session database

Use the supported `show` and `clear` commands to manage and display information about the bulk leasequery operation for the DHCP relay agent and the DHCP local server. See ["Verifying and Managing DHCP Individual and Bulk Leasequery Configurations" on page 1157](#).

EVPN-MPLS DHCPv6-PD Stateful Relay Synchronization for Active-Active mode (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024 and ACX7024X)

This DHCPv6-PD stateful relay synchronization feature provides support for DHCPv6 prefix delegation configuration, that includes:

- Synchronization between two leaf nodes of EVPN-MPLS for active-active mode
- DHCPv6 prefix delegation to automate the delegation of IPv6 prefixes to the CPE.
- Support for bulk leasequery on a DHCPv6 relay agent.
- Configuring parameters that the DHCP relay agent uses when sending DHCP bulk leasequery messages to obtain lease information from the DHCP local servers in the logical system/routing instance.
- EVPN-VXLAN DHCPv6 state synchronization using Active-lease-query/Bulk-lease-query over IRB.

[See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#), [DHCP Leasequery Methods](#), [active-leasequery \(DHCP Relay Agent\)](#) and [bulk-leasequery \(DHCP Relay Agent\)](#) .]

Configuration Example:

The following example displays an Active-Active configuration with stale-timer setting. Stale-timer configuration is required to support an Active-Active lease query. This configuration optimizes the synchronization time when both the peers get the solicit packets at the same time.

```

dhcp-relay
{
    dhcpv6
    {
        group v6relay
        {
            active-server-group v6server;
            interface irb.0;
        }
        relay-agent-interface-id
        {
            include-irb-and-l2;
        }
        server-group
        {
            v6server
            {
                1000::1;
            }
        }
        bulk-leasequery;
        active-leasequery
        {
            peer-address
            {
                1003::1;
            }
        }
    }
    overrides
    {
        always-write-option-82;
    }
    relay-option-82
    {
        circuit-id
    }
}

```

```

        {
            include-irb-and-l2;
        }
    }

    server-group {
        v4server {
            100.0.0.1;
        }
    }

    group v4relay {
        active-server-group v4server;
        interface irb.0;
    }

    stale-timer 20;
    bulk-leasequery;
    active-leasequery {
        peer-address {
            103.0.0.1;
        }
    }
}

dhcp-relay {
    dhcpv6 {
        group v6relay {
            active-server-group v6server;
            interface irb.0;
        }

        relay-agent-interface-id {
            include-irb-and-l2;
        }

        server-group {
            v6server {
                1000::1;
            }
        }

        bulk-leasequery;
        active-leasequery {
            peer-address {
                1002::1;
            }
        }
    }

    overrides {

```

```

always-write-option-82;
}
relay-option-82 {
circuit-id {
include-irb-and-l2;
}
}
server-group {
v4server {
100.0.0.1;
}
}
group v4relay {
active-server-group v4server;
interface irb.0;
}
stale-timer 20;
bulk-leasequery;
active-leasequery {
peer-address {
102.0.0.1;
}
}
}

```

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database

You must issue a request command to trigger the DHCP relay agent to initiate an individual leasequery or bulk leasequery operation, which requests current lease information from DHCP local servers. Each individual leasequery updates the DHCP relay agent's lease database with information for an individual client. Each bulk leasequery updates the relay agent's lease database for multiple clients. [Table 48 on page 1152](#) lists the various query options that are available for DHCPv4, DHCPv6, individual leasequery, and bulk leasequery.

Table 48: Query Options for Each Leasequery Method

Query Option	DHCPv4 Individual Leasequery	DHCPv4 Bulk Leasequery	DHCPv6 Individual Leasequery	DHCPv6 Bulk Leasequery
Agent Remote ID	-	✓	-	✓

Table 48: Query Options for Each Leasequery Method (*Continued*)

Query Option	DHCPv4 Individual Leasequery	DHCPv4 Bulk Leasequery	DHCPv6 Individual Leasequery	DHCPv6 Bulk Leasequery
Client ID	✓	✓	–	–
Client ID (DUID)	–	–	✓	✓
Gateway Address	✓ mandatory	–	–	–
IPv4 Address	✓	✓	–	–
IPv6 Prefix	–	–	✓	✓
Link Address	–	–	–	✓
MAC Address	✓	✓	–	–
Relay Agent ID	–	✓	–	✓



NOTE: When you have configured DHCPv6 bulk leasequery on a relay agent with the `bulk-leasequery` statement and the `trigger automatic` option, you do not initiate the query with a `request` command. Instead, the query is automatically triggered whenever the `jdhcpd` process on the relay agent starts (for example, after a `jdhcpd` restart, a relay agent device reboot, a *graceful Routing Engine switchover*, or a unified ISSU) and there are no bound subscribers in the session database. The automatic bulk leasequery is always based on the relay agent Relay-ID option (option 53).

When the automatic trigger support is configured, you can still use the `request` command to manually trigger bulk leasequeries separate from the automatic queries.



NOTE: Active leasequery does not require a `request` command for initiation. Instead, it is automatically initiated when you configure it. Active leasequery does require you to configure bulk leasequery.

DHCPv4 relay agents can have multiple interfaces with different IP addresses, so that each interface can act as a gateway for different set of clients. This means that you must always specify the gateway address in your request.

To initiate a DHCPv4 individual leasequery to update binding information, you must always specify the gateway IP address of the relay agent. You must also specify the type of query:

- Specify an IP address leased to the client.

```
user@host> request dhcp relay leasequery ipv4-address gateway-address giaddr
```

- Specify the client's MAC address.

```
user@host> request dhcp relay leasequery mac-address gateway-address giaddr
```

- Specify the client identifier (option 61).

```
user@host> request dhcp relay leasequery client-id gateway-address giaddr
```

To initiate a DHCPv4 bulk leasequery to update binding information, you can:

- Specify an IP address leased to the client.

```
user@host> request dhcp relay bulk-leasequery ipv4-address
```

- Specify the client's MAC address.

```
user@host> request dhcp relay bulk-leasequery mac-address
```

- Specify the client identifier option (option 61).

```
user@host> request dhcp relay bulk-leasequery client-id
```

- Specify the Relay Agent Identifier suboption (suboption 12) of the DHCP relay agent information option (option 82).

```
user@host> request dhcpv6 relay bulk-leasequery relay-id relay-id
```

By default, the bulk leasequery operation uses the relay ID of the DHCPv4 relay agent if you do not explicitly specify any of the following options: *client-id*, *ipv4-address*, *mac-address*, *relay-id*, or *remote-id*.

```
user@host> request dhcpv6 relay bulk-leasequery
```

- Specify the Agent Remote ID (suboption 2) of the DHCPv4 relay agent information option (option 82).

```
user@host> request dhcpv6 relay bulk-leasequery remote-id remote-id
```

To initiate a DHCPv6 individual leasequery to update binding information, you can:

- Specify the client ID (option 1).

```
user@host> request dhcpv6 relay leasequery client-id
```

- Specify an IPv6 address leased to the client.

```
user@host> request dhcpv6 relay leasequery ipv6-prefix
```

To initiate a DHCPv6 bulk leasequery to update binding information, you can:

- Specify the client ID (option 1).

```
user@host> request dhcpv6 relay bulk-leasequery client-id
```

- Specify the IPv6 prefix.

```
user@host> request dhcpv6 relay bulk-leasequery ipv6-prefix
```

- Specify the IPv6 link address.

```
user@host> request dhcpv6 relay bulk-leasequery link-address ipv6-link-address
```

- Specify the Relay-ID option (option 53).

```
user@host> request dhcpv6 relay bulk-leasequery relay-id relay-id
```

By default, the bulk leasequery operation uses the relay ID of the DHCPv6 relay agent if you do not explicitly specify any of the following options: *client-id*, *ipv6-prefix*, *ipv6-link-address*, *relay-id*, or *remote-id*.

```
user@host> request dhcpv6 relay bulk-leasequery
```

- Specify the Relay Agent Remote-ID option (option 37).

```
user@host> request dhcpv6 relay bulk-leasequery remote-id remote-id
```

For any individual and bulk leasequery request, in addition to the options listed above, you can optionally specify qualifiers to limit the query to particular DHCP servers. Otherwise the query is sent to all DHCP servers known to the relay agent.

You can specify an address for the local server or the name of a group of local servers. You can specify a logical system, a routing-instance, or both, either alone or in addition to the server address or group.



NOTE: In the following example, *option* means any configurable option as shown earlier. For brevity, the example shows only a DHCPv4 individual leasequery and only some of the possibilities. For more information, see the individual command topics: *request dhcp relay leasequery*, *request dhcpv6 relay leasequery*, *request dhcp relay bulk-leasequery*, and *request dhcpv6 relay bulk-leasequery*.

- Specify an address for the local server.

```
user@host> request dhcp relay leasequery option server-address address
```

- Specify a logical system.

```
user@host> request dhcp relay leasequery option logical-system logical-system-name
```

- Specify a routing instance and a named group of local servers.

```
user@host> request dhcp relay leasequery option routing-instance routing-instance-name server-  
group group-name
```

Verifying and Managing DHCP Individual and Bulk Leasequery Configurations

IN THIS SECTION

- Purpose | 1157
- Action | 1157

Purpose

View or clear information about DHCP individual leasequery and bulk leasequery operations. Use the supported `show` and `clear` commands to manage and display information about the leasequery and bulk leasequery operations; for the DHCP relay agent and the DHCP local server.



NOTE: For active leasequery, see "[Verifying and Managing DHCP Active Leasequery Operations](#)" on page 1158.

Action

Use the supported `show` and `clear` commands to manage and display information about the leasequery operations for the DHCP relay agent and the DHCP local server.

- To display leasequery information for DHCPv4 or DHCPv6 relay agent:

```
user@host> show dhcp relay statistics (leasequery | bulk-leasequery-connections)  
user@host> show dhcpv6 relay statistics (leasequery | bulk-leasequery-connections)
```

- To clear leasequery information for DHCPv4 or DHCPv6 relay agent:

```
user@host> clear dhcp relay statistics (leasequery | bulk-leasequery-connections)  
user@host> clear dhcpv6 relay statistics (leasequery | bulk-leasequery-connections)
```

- To display leasequery information for DHCPv4 or DHCPv6 local server:

```
user@host> show dhcp server statistics bulk-leasequery-connections
user@host> show dhcpv6 server statistics bulk-leasequery-connections
```

- To clear leasequery information for DHCPv4 or DHCPv6 local server:

```
user@host> clear dhcp server statistics bulk-leasequery-connections
user@host> clear dhcpv6 server statistics bulk-leasequery-connections
```

Verifying and Managing DHCP Active Leasequery Operations

IN THIS SECTION

- Purpose | 1158
- Action | 1158

Purpose

View or clear information about DHCP active leasequery operations. Use the supported `show` and `clear` commands to manage and display information about the active leasequery operations; for the DHCP relay agent and the DHCP local server.



NOTE: For DHCP individual and bulk leasequery, see "[Verifying and Managing DHCP Individual and Bulk Leasequery Configurations](#)" on page 1157.

Action

Use the supported `show` and `clear` commands to manage and display information about the leasequery operations for the DHCP relay agent and the DHCP local server.

- To display active leasequery information for DHCPv4 or DHCPv6 peer relay agents:

```
user@host> show dhcp relay active-leasequery
user@host> show dhcpv6 relay active-leasequery
```

- To clear active leasequery information for DHCPv4 or DHCPv6 relay agent:

```
user@host> clear dhcp relay active-leasequery statistics
user@host> clear dhcpv6 relay active-leasequery statistics
```

- To display information about active leasequery neighbors:

```
user@host> show dhcp active-leasequery neighbors
user@host> show dhcpv6 active-leasequery neighbors
```

You can display general information for all peers. You can also display statistics for specific peers and specific access interfaces. For example:

- For each pseudowire interface on the BNG, display the IP address of the BNG neighbor associated with the interface.

```
user@host> show dhcp active-leasequery neighbors
```

Interface	Neighbor Address
ps2.0	198.51.100.5
ps1.0	198.51.100.7

- Display statistics for DHCPv4 and DHCPv6 peers.

```
user@host> show dhcp relay active-leasequery statistics peer 198.51.100.1
```

```

peer : 198.51.100.1
Topology-Discover Configured      : Yes
State                             : Done
Bindings Sent                     : 0
Bindings Received                 : 0
Bindings Installed Successfully   : 0
Bindings Failed to install       : 0
Last Synchronization Time        : None
ALQ Transmit Buffer count         : 0x ffff
Max Leasequery Transmit Rate     : 60
```

```

Local Interface count      : 2
Remote Interface count    : 2

```

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::2
```

```

peer : 2001:db8::2
Topology-Discover Configured      : Yes
State                             : Done
Bindings Sent                     : 8112
Bindings Received                 : 12382
Bindings Installed Successfully   : 0
Bindings Failed to install       : 0
Last Synchronization Time        : 2020-02-05 01:27:54 IST
ALQ Transmit Buffer count         : 0x ffff
Max Leasequery Transmit Rate     : 60
Local Interface count            : 2
Remote Interface count           : 2

```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
20.1R1	(Support for ps interfaces added in Junos OS Release 20.1R1.)
19.2R1	Starting in Junos OS Release 19.2R1, topology discovery enables DHCP relay peers to discover information about each other's subscriber interfaces.
19.1R1	Starting in Junos OS Release 19.1R1, DHCP active leasequery addresses the situation where it is desirable for the relay agent to receive periodic updates of client information to keep up with dynamic DHCP binding activity.
16.1	Starting in Junos OS Release 16.1, subscriber management supports the individual leasequery feature, which enables the DHCPv4 or DHCPv6 relay agent to quickly and efficiently obtain the current lease information from a DHCP local server.
16.1	Starting in Junos OS Release 16.1, subscriber management supports the bulk leasequery feature, which enables each request from the DHCP relay agent to retrieve lease information for multiple subscribers in bulk from a configured DHCP server in a programmed manner.

RELATED DOCUMENTATION

DHCP Overview 1015
DHCPv6 Local Server 1449
DHCPv6 Relay Agent 1455

DHCP Client Authentication With An External AAA Authentication Service

IN THIS SECTION

- [Specifying Authentication Support | 1161](#)
- [Creating Unique Usernames for DHCP Clients | 1162](#)
- [Example-Configuring DHCP with External Authentication Server | 1164](#)

Specifying Authentication Support

Include the `authentication` statement at hierarchy levels given in [Table 49 on page 1161](#). You can configure either global authentication support or group-specific support.

Table 49: Supported Hierarchy Levels for Authentication Support

Supported Hierarchy Level	Hierarchy Level
DHCP local server	[edit system services dhcp-local-server]
DHCP relay agent	[edit forwarding-options dhcp-relay]
DHCPv6 local server	[edit system services dhcp-local-server dhcpv6]
DHCPv6 relay agent	[edit forwarding-options dhcp-relay dhcpv6]

Creating Unique Usernames for DHCP Clients

You can configure the extended DHCP application to include additional information in the username that is passed to the external AAA authentication service when the DHCP client logs in. This additional information enables you to construct usernames that uniquely identify subscribers (DHCP clients).

To configure unique usernames, use the `username-include` statement. You can include any or all of the additional statements.

```
authentication {
  username-include {
    circuit-type;
    client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (DHCP Local Server) (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;
  }
}
```



NOTE: If you do not include a username in the authentication configuration, the router (or switch) does not perform authentication; however, the IP address is provided by the local pool if it is configured.

When you use the DHCPv6 local server, you must configure authentication and the client username; otherwise client login fails.

The following list describes the optional information that you can include as part of the username:

- `circuit-type`—The circuit type used by the DHCP client, for example `enet`.
- `client-id`—The client identifier option (option 1). (DHCPv6 local server DHCPv6 relay agent only)
- `delimiter`—The delimiter character that separates components that make up the concatenated username. The default delimiter is a period (.). The semicolon (;) is not supported as a delimiter character.
- `domain-name`—The client domain name as a string. The router adds the @ delimiter to the username.

- `interface-description`—The description of the device (physical) interface or the logical interface.
- `interface-name`—The interface name, including the interface device and associated VLAN IDs.
- `logical-system-name`—The name of the logical system, if the receiving interface is in a logical system.
- `mac-address`—The client MAC address, in a string of the format `xxxx.xxxx.xxxx`.
- `option-60`—The portion of the option 60 payload that follows the length field. (Not supported for DHCPv6 local server)
- `option-82 <circuit-id> <remote-id>`—The specified contents of the option 82 payload. (Not supported for DHCPv6 local server)
 - `circuit-id`—The payload of the Agent Circuit ID suboption.
 - `remote-id`—The payload of the Agent Remote ID suboption.
 - Both `circuit-id` and `remote-id`—The payloads of both suboptions, in the format: `circuit-id[delimiter]remote-id`.
 - Neither `circuit-id` or `remote-id`—The raw payload of the option 82 from the PDU is concatenated to the username.



NOTE: For DHCP relay agent, the option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.

- `relay-agent-interface-id`—The Interface-ID option (option 18). (DHCPv6 local server or DHCPv6 relay agent only)
- `relay-agent-remote-id`—The DHCPv6 Relay Agent Remote-ID option (option 37). (DHCPv6 local server or DHCPv6 relay agent only)
- `relay-agent-subscriber-id`—(On routers only) The DHCPv6 Relay Agent Subscriber-ID option (option 38). (DHCPv6 local server or DHCPv6 relay agent only)
- `routing-instance-name`—The name of the routing instance, if the receiving interface is in a routing instance.
- `user-prefix`—A string indicating the user prefix.
- `vlan-tags`—The subscriber VLAN tags. Includes the outer VLAN tag and, if present, the inner VLAN tag. You can use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

For DHCPv6 clients, because the DHCPv6 packet format has no specific field for the client MAC address, the MAC address is derived from among several sources with the following priority:

- Client DUID Type 1 or Type 3.
- Option 79 (client link-layer address), if present.
- The packet source address if the client is directly connected.
- The link local address.

The router (switch) creates the unique username by including the specified additional information in the following order, with the fields separated by a delimiter.

For DHCP local server and DHCP relay agent:

```
user-prefix[delimiter]mac-address[delimiter]logical-system-name[delimiter]routing-instance-
name[delimiter]circuit-type[delimiter]interface-
name[delimiter]option-82[delimiter]option-60@domain-name
```

For DHCPv6 local server:

```
user-prefix[delimiter]mac-address[delimiter]logical-system-name[delimiter]routing-instance-
name[delimiter]circuit-type[delimiter]interface-name[delimiter]relay-agent-remote-
id[delimiter]relay-agent-subscriber-id[delimiter]relay-agent-interface-id[delimiter]client-
id@domain-name
```

Example-Configuring DHCP with External Authentication Server

To configure authentication at DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent levels.

1. Specify that you want to configure authentication.

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

2. (Optional) Specify the optional information you want to include in the username.

```
[edit system services dhcp-local-server authentication username-include]
user@host# set username-include circuit-type
user@host# set username-include domain-name example.com
```

```
user@host# set username-include mac-address
user@host# set username-include user-prefix wallybrown
```

3. Configure an optional password that the extended DHCP application presents to the external AAA authentication service to authenticate the specified username.

```
[edit system services dhcp-local-server authentication]
user@host# set password $ABC123
```

The following example shows a sample configuration that creates a unique username. The username is shown after the configuration.

```
authentication {
  username-include {
    circuit-type;
    domain-name example.com;
    mac-address 2001:db8::/32;
    user-prefix wallybrown;
  }
}
```

The resulting unique username is:

```
wallybrown.2001:db8::/32.enet@example.com
```

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

Receiving DHCP Options From a RADIUS Server

IN THIS SECTION

- [Centrally Configure DHCP Options on a RADIUS Server | 1166](#)
- [Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview | 1170](#)
- [Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6 Subscribers | 1174](#)
- [Monitoring DHCP Options Configured on RADIUS Servers | 1176](#)

Centrally Configure DHCP Options on a RADIUS Server

IN THIS SECTION

- [RADIUS-Sourced Options | 1166](#)
- [Client-Sourced Options Configuration | 1167](#)
- [Data Flow for RADIUS-Sourced DHCP Options | 1167](#)
- [Multiple VSA 26-55 Instances Configuration | 1169](#)
- [DHCP Options That Cannot Be Centrally Configured | 1169](#)

DHCP management on Junos OS devices support central configuration of DHCP options directly on the RADIUS server (RADIUS-sourced options) and traditional client-sourced options configuration. Read the following sections for information on central configuration of DHCP options on the RADIUS server.

RADIUS-Sourced Options

Subscriber management (on the routers) or DHCP management (on the switches) enables you to centrally configure DHCP options on a RADIUS server and then distribute the options on a per-subscriber or per DHCP-client basis. This method results in RADIUS-sourced DHCP options—the DHCP options originate at the RADIUS server and are sent to the subscriber (or DHCP client). This differs from the traditional client-sourced method (also called DHCP-sourced) of configuring DHCP options, in which the options originate at the client and are sent to the RADIUS server. The subscriber management (DHCP management) RADIUS-sourced DHCP options are also considered to be *opaque*, because DHCP local server performs minimal processing and error checking for the DHCP options string before passing the options to the subscriber (DHCP client).

Subscriber management (or DHCP management) uses Juniper Networks VSA 26-55 (DHCP-Options) to distribute the RADIUS-sourced DHCP options. The RADIUS server includes VSA 26-55 in the Access-Accept message that the server returns during subscriber authentication or DHCP client authentication. The RADIUS server sends the Access-Accept message to the RADIUS client, and then on to DHCP local server for return to the DHCP subscriber. The RADIUS server can include multiple instances of VSA 26-55 in a single Access-Accept message. The RADIUS client concatenates the multiple instances and uses the result as a single instance.

There is no CLI configuration required to enable subscriber management (DHCP management) to use the centrally configured DHCP options—the procedure is triggered by the presence of VSA 26-55 in the RADIUS Access-Accept message.

When building the offer packet for the DHCP client, DHCP local server uses the following sequence:

1. Processes any RADIUS-configured parameters that are passed as separate RADIUS attributes; for example, RADIUS attribute 27 (Session Timeout).
2. Processes any client-sourced parameters; for example, RADIUS attributes 53 (DHCP Message Type) and 54 (Server Identifier).
3. Appends (without performing any processing) the opaque DHCP options string contained in the VSA 26-55 received from the RADIUS server.

Client-Sourced Options Configuration

In addition to supporting central configuration of DHCP options directly on the RADIUS server (RADIUS-sourced options), subscriber management (DHCP management) also supports the traditional client-sourced options configuration, in which the router's (switch's) DHCP component sends the options to the RADIUS server. The client-sourced DHCP options method is supported for both DHCP local server and DHCP relay agent; however, the RADIUS-sourced central configuration method is supported on DHCP local server only. Both the RADIUS-sourced and client-sourced methods support DHCPv4 and DHCPv6 subscribers (clients).




NOTE: You can use the RADIUS-sourced and client-sourced methods simultaneously on DHCP local server. However, you must ensure that the central configuration method does not include options that override client-sourced DHCP options, because this can create unpredictable results.

Data Flow for RADIUS-Sourced DHCP Options

Figure 7 on page 1168 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

Figure 7: DHCP Options Data Flow



The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

1. The subscriber (DHCP client) sends a DHCP discover message (or DHCPv6 solicit message) to the DHCP local server. The message includes client-sourced DHCP options.
2. The DHCP local server initiates authentication with the Junos OS RADIUS client.
3. The RADIUS client sends an Access-Request message on behalf of the subscriber (DHCP client) to the external RADIUS server. The message includes the subscriber's (DHCP client's) client-sourced DHCP options.
4. The external RADIUS server responds by sending an Access-Accept message to the RADIUS client. The Access-Accept message includes the RADIUS-sourced opaque DHCP options in VSA 26-55.
5. The RADIUS client sends the DHCP options string to DHCP local server. If there are multiple VSA 26-55 instances, the RADIUS client first assembles them into a single options string.
6. DHCP local server processes all options into the DHCP offer (or DHCPv6 reply) message, except for the RADIUS-sourced VSA 26-55 DHCP options. After processing all other options, DHCP local server then appends the unmodified VSA 26-55 DHCP options to the message and sends the message to the subscriber (DHCP client).
7. The subscriber (DHCP client) is configured with the DHCP options.
8. The following operations occur after the subscriber (DHCP client) receives the DHCP options:

- **Accounting**—The RADIUS client sends Acct-Start and Interim-Accounting requests to the RADIUS server, including the RADIUS-sourced DHCP options in VSA 26-55. By default, the DHCP options are included in accounting requests.
- **Renewal**—When the subscriber (DHCP client) renews, the cached DHCP options value is returned in the DHCP renew (or DHCPv6 ACK) message. The originally assigned DHCP options cannot be modified during a renew cycle.
- **Logout**—When the subscriber (DHCP client) logs out, the RADIUS client sends an Acct-Stop message to the RADIUS server, including the RADIUS-sourced VSA 26-55.

Multiple VSA 26-55 Instances Configuration

VSA 26-55 supports a maximum size of 247 bytes. If your RADIUS-sourced DHCP options field is greater than 247 bytes, you must break the field up and manually configure multiple instances of VSA 26-55 for the RADIUS server to return. When using multiple instances for an options field, you must place the instances in the packet in the order in which the fragments are to be reassembled by the RADIUS client. The fragments can be of any size of 247 bytes or less.



BEST PRACTICE: For ease of configuration and management of your DHCP options, you might want to have one DHCP option per VSA 26-55 instance, regardless of the size of the option field.

When the RADIUS client returns a reassembled opaque options field in an accounting request to the RADIUS server, the client uses 247-byte fragments. If you had originally created instances of fewer than 247 bytes, the returned fragments might not be the same as you originally configured on the RADIUS server.



NOTE: If you are configuring Steel-Belted Radius (SBR) to support multiple VSA 26-55 instances, ensure that you specify VSA 26-55 with the R0 flags in the Subscriber Management RADIUS dictionary file. The R value indicates a multivalued reply attribute and the 0 value indicates an ordered attribute.

DHCP Options That Cannot Be Centrally Configured

Table 50 on page 1170 shows the DHCP options that you must not centrally configure on the RADIUS server.

Table 50: Unsupported Opaque DHCP Options

DHCP Option	Option Name	Comments
Option 0	Pad Option	Not supported.
Option 51	IP Address Lease Time	Value is provided by RADIUS attribute 27 (Session-Timeout).
Option 52	Option Overload	Not supported.
Option 53	DHCP Message Type	Value is provided by DHCP local server.
Option 54	Server Identifier	Value is provided by DHCP local server.
Option 55	Parameter Request List	Value is provided by DHCP local server.
Option 255	End	Value is provided by DHCP local server.
-	DHCP magic cookie	Not supported.

SEE ALSO
[DHCP with External Authentication Server](#)
[DHCP Overview](#)
[IP Address Assignment Pool](#)
Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview**IN THIS SECTION**

- [Differentiating Subscriber Classes with DHCPv6 Option 17 and VSA 26-207 | 1172](#)
- [Excluding the VSAs from RADIUS Messages | 1173](#)

The RADIUS server, which is configured independently of DHCPv4 and DHCPv6, authenticates clients and supplies the IPv4 or IPv6 prefix and client configuration parameters. To establish the client sessions on the network, the DHCPv4 and DHCPv6 parameters are sent from the client device through the DHCP (either DHCPv4 or DHCPv6) server to the RADIUS server and vice versa. Starting in Junos OS Release 17.4R1, the exchange of parameters is enhanced with the introduction of several new vendor-specific attributes (VSAs) and changes to the existing DHCP-Options VSA (26-55).

An immediate interim accounting report is sent to the RADIUS server when configurable events occur, such as a change in family state. When these events occur, the RADIUS server has no direct way to determine the reason for the report. You can use the Acct-Request-Reason VSA (26-210) to send the reason in the start accounting report as well as in the immediate interim accounting report.

The broadband network gateway (BNG) sends an interim accounting report to the RADIUS server whenever the second family (either IPv4 or IPv6) of a dual-stack session (DHCPv4, DHCPv6, or PPPoE) is activated or the first family (either IPv4 or IPv6) of a dual-stack session (DHCPv4, DHCPv6, or PPPoE) is deactivated. For the immediate interim accounting report to be sent, configure the `family-state-change-immediate-update` statement on the BNG at the `[edit access profile profile-name accounting]` hierarchy level.

The following VSAs are used for exchanging the client parameters with the RADIUS server:

- DHCPv6-Options VSA (26-207):
 - The DHCPv6-Options VSA (26-207) is used to exchange DHCPv6 options with the RADIUS server. In releases earlier than Junos OS Release 17.4R1, the DHCPv6 options are included with DHCPv4 options in the DHCP-Options VSA (26-55).

The option values sent from the DHCPv6 client to the DHCPv6 server are saved in the session database separately from the values sent from the DHCPv6 server to the DHCPv6 client.

 - If the DHCPv6 options are too large to fit in one VSA, then they are split into multiple, sequential VSAs in the RADIUS packet. In this case, the options are split at the VSA size limit rather than at the type-length-value (TLV) boundary.
 - If multiple instances of the VSA are included in the RADIUS Access-Accept message, then they are concatenated into a single block and stored in the session database without checking the TLV for validity.
- DHCP-Options VSA (26-55):
 - The DHCP-Options VSA (26-55) is used to exchange DHCPv4 options with the RADIUS server.
 - With the introduction of VSA 26-207, VSA 26-55 includes only DHCPv4 options.
 - If the DHCPv4 options are too large to fit in one VSA, then they are split into multiple, sequential VSAs in the RADIUS packet. In this case, the options are split at the VSA size limit rather than at the TLV boundary.

- If multiple instances of the VSA are included in the RADIUS Access-Accept message, then they are concatenated into a single block and stored in the session database without checking the TLV for validity.
- DHCP-Header VSA (26-208):
 - The DHCP-Header VSA (26-208) conveys the DHCPv4 packet header to the RADIUS server. The header information is used for instantiating dynamic subscriber interfaces.
 - The VSA is allowed only in RADIUS Access-Request messages and is stored in the session database.
- DHCPv6-Header VSA (26-209):
 - The DHCPv6-Header VSA (26-209) conveys the DHCPv6 packet header to the RADIUS server. The header information is used for instantiating dynamic subscriber interfaces.
 - The VSA is allowed only in RADIUS Access-Request messages and is stored in the session database.
- Acct-Request-Reason VSA (26-210):
 - The Acct-Request-Reason VSA (26-210) conveys the reason for sending an accounting request. The VSA is included only in RADIUS Acct-Start and Interim-Update messages. The VSA is present only for subscriber accounting reports; it is not present for service session or Extensible Subscriber Services Manager (ESSM) reports.
 - The typical value for the VSA in Acct-Start messages is IP active (0x0004) or IPv6 active (0x0010), indicating that the IPv4 or IPv6 address family has been activated. For Layer 2 wholesale VLAN networks, the value is Session active (0x0040), because there is no IPv4 or IPv6 family. The value for MLPPP is also Session active, because accounting messages are sent for the link session rather than the bundle session. ESSM sessions are child sessions of a parent subscriber session and are treated as ESSM service sessions. The VSA is sent only for the parent subscriber session.

Differentiating Subscriber Classes with DHCPv6 Option 17 and VSA 26-207

Starting in Junos OS Release 18.3R1, you can use the DHCPv6-Options VSA (26-207) to differentiate between different classes of subscribers during DHCPv6 relay authentication. For example, you may want to assign different IPv6 prefixes to different subscriber classes.

You must configure your RADIUS server to include the following information in the VSA:

- Juniper Networks enterprise number, 2636
- Suboption 5, JDHCPD_VS_OPT_CODE_KT_SUBSCRIBER_CLASS



NOTE: To configure this information, refer to the documentation for your RADIUS server. You must encode the information in the DHCPv6 options format in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*.

You set a different value for suboption 5 for each class you want to differentiate. You develop your own scheme to determine the mapping between value and class.

VSA 26-207 conveys the subscriber class information in the Access-Accept message returned by the RADIUS server during DHCPv6 subscriber authentication. The contents of the VSA are passed from the AAA process to the DHCP process in the session database attribute, SDB_SERVER_DHCPV6_OPTIONS. The DHCPv6 relay agent extracts the information from the SDB attribute and places it in DHCPv6 option 17. The relay agent subsequently passes option 17 to the DHCPv6 local server in the Relay-Forward header. The local server can then return the relay agent configuration and service information specific to the relevant subscriber classes.

In releases earlier than Junos OS 18.3R1, only the DHCP local server supports VSA 26-207. Only suboption 1 (JDHCPD_VS_OPT_CODE_HOST_NAME) and suboption 4 (JDHCPD_VS_OPT_CODE_LOCATION_NAME) are supported. The DHCP relay agent also discards the SDB_SERVER_DHCPV6_OPTIONS attribute if it is received.

Suboptions received from RADIUS have a higher precedence than the information configured locally. For example, if the host name and the location are configured with the `host-name` statement at the `[edit forwarding-options dhcp-relay dhcpv6 relay-option-vendor-specific]` hierarchy level and they are received in suboptions 1 and 4 from RADIUS, the RADIUS values are used.

Excluding the VSAs from RADIUS Messages

You can exclude any of these VSAs from being sent by using the `exclude` statement as shown in the following example:

```
[edit access profile profile-name radius attributes]
user@host# set exclude acct-request-reason [accounting-start | accounting-stop]
user@host# set exclude dhcp-header [access-request]
user@host# set exclude dhcpv6-header [access-request]
user@host# set exclude dhcpv6-options [access-request | accounting-start | accounting-stop]
```

Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6 Subscribers

IN THIS SECTION

- [Client Options | 1174](#)
- [Exchange of DHCPv4 Client, DHCPv4 Server, and RADIUS-Sourced Options | 1175](#)
- [Exchange of DHCPv6 Client, DHCPv6 Server, and RADIUS-Sourced Options | 1175](#)

The Dynamic Host Configuration Protocol (DHCP) server can serve as a DHCP local server, a DHCP client, or a DHCP relay agent, for both DHCPv4 and DHCPv6 subscribers.

Currently, some of the client parameters—for example, the DHCPv4 and DHCPv6 packet header—cannot be passed to and from the RADIUS server. From Junos OS Release 17.4 onward, enhancements are made to facilitate better communication between the DHCP servers (both DHCPv4 and DHCPv6) and the RADIUS server. The client parameters are saved in a session database and sent to the RADIUS server; and the RADIUS server, in turn, authenticates the client and also responds with the options to be sent back to that client.

Client Options

The client options can be configured in multiple locations such as DHCPv4 or DHCPv6 servers, or in the RADIUS server. If the client configuration is available in multiple locations, a conflict can arise regarding the source of the configuration details. In case of such conflicts, the following order of preference is considered:

- Options received from the RADIUS server through vendor-specific attributes (VSAs)
- Options received from the RADIUS server through the respective session databases
- Options from the DHCP local configuration, which are present on the DHCP server

As an example of the aforementioned preference, consider the case of DHCPv4 lease time. If the `AUTHD_ATTR_SESSION_TIMEOUT` option, which is a VSA stored in the RADIUS server, is returned from the RADIUS server, preference is given to it. If this option is not returned, preference is given to option 51 in respective session database for DHCPv4. If that option is also not returned, the option is sourced from DHCP local configuration.

Similarly, for DHCPv6 lease time, the first preference is given to the `AUTHD_ATTR_SESSION_TIMEOUT` VSA from the RADIUS server. If `AUTHD_ATTR_SESSION_TIMEOUT` is not present, the RADIUS-sourced option `valid-lifetime`

or preferred-lifetime takes the precedence. If that is also not available, then the option is sourced from the DHCPv6 local configuration.

Exchange of DHCPv4 Client, DHCPv4 Server, and RADIUS-Sourced Options

The following steps illustrate the process of exchange of configuration options between a DHCPv4 client, a DHCPv4 server, and the RADIUS server:

- A *discover* message from a DHCPv4 client is received by the DHCPv4 server.
- The DHCP option is saved to the respective session database.

In Junos OS releases before 17.4R1, the same attribute is used to store both DHCPv4 and DHCPv6 options. However, with the support for single-session DHCP dual-stack, there are separate session database attributes for DHCPv4 and DHCPv6.

- The DHCP header information is saved in the session database.

A new session database attribute is added to store the header information, and this information is sent to the RADIUS server for authentication.

- An *access request* message is sent from the DHCPv4 server to the RADIUS server, and when an *access accept* message is received from the RADIUS server, the DHCPv4 options are saved to the respective session database attributes and sent to the client.
- DHCPv4 server-specific options are added to the packet.



NOTE: The DHCPv4 server can source both solicited and unsolicited options from the local configuration. Thus, it is important to prevent duplication while the options are added.

- DHCPv4 lease information is extracted from the RADIUS-sourced DHCP option 51.

The respective session database attribute is used to check whether option 51 (lease time) is sourced by RADIUS. If it is, then the attribute value is extracted and saved in the client data structure. If it is not sourced by RADIUS, the attribute value is taken from the local pool configuration or the DHCPv4 attribute configuration, which is an existing functionality. A similar check is performed for option 58 (renewal time (T1)) and option 59 (rebinding time (T2)).

- An *offer message* is sent from the DHCPv4 server to DHCPv4 client.

Exchange of DHCPv6 Client, DHCPv6 Server, and RADIUS-Sourced Options

The following steps illustrate the process of exchange of configuration options between a DHCPv6 client, a DHCPv6 server, and the RADIUS server:

- A *solicit* message from a DHCP client is received by the DHCPv6 server.
- DHCPv6 options are saved in the session database of the DHCPv6 server.

In Junos OS releases before 17.4R1, DHCPv6 options are saved in the respective session database attribute. Because of the current single-session DHCP dual-stack support, there is need to have separate session database attributes for saving DHCPv4 and DHCPv6 options. If the client is part of a single-session dual-stack configuration, the respective DHCPv6 options session database attribute is used. The DHCPv6 options are directly copied to the session database without any changes and then sent to the RADIUS server.



NOTE: DHCPv6 auth-option (option 11) is also part of these options.

- A DHCPv6 message header is saved to the session database.

A new session database attribute is added to copy the DHCPv6 message header.

- An *access request* message is sent from the DHCPv6 server, which in turn receives an *access accept* message from the RADIUS server. This message contains RADIUS-sourced DHCPv6 options that are stored in a new session database attribute.
- DHCPv6 lease information is extracted from the RADIUS-sourced DHCPv6 option.

In case of DHCPv6, the lease time is embedded within the options `OPTION_IA_NA` and `OPTION_IA_PD`. Client lease time starts with these values from the RADIUS Server. If the `IA_ADDRESS`, `IA_PREFIX`, `IA_NA`, or `IA_PD` option is not sourced from RADIUS, then these options are taken from the local pool and delegated pool configuration.

- DHCPv6 server-specific options are added to the packet.



NOTE: A DHCPv6 server can source both solicited and unsolicited options from the local configuration. Thus, it is important to prevent duplication while the options are added.

- An *advertise* message is sent from the DHCPv6 server to the DHCPv6 client.

Monitoring DHCP Options Configured on RADIUS Servers

IN THIS SECTION

 Purpose | 1177

- [Action | 1177](#)
- [Meaning | 1177](#)

Purpose

View information for DHCP options that are centrally configured on a RADIUS server and that are distributed using Juniper Networks VSA 26-55 (DHCP-Options).

Action

To display information for opaque DHCP options:

```
user@host> show subscribers detailType: DHCP
IP Address: 192.168.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-prof-23
MAC Address: 00:00:5E:00:53:98
State: Active
Radius Accounting ID: jnpr :2304
Session Timeout (seconds): 3600
Idle Timeout (seconds): 600
Login Time: 2011-08-25 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
```

Meaning

```
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
```

The DHCP options output provides the following information:

- The len field is the total number of hex values in the message.
- The hex values specify the type, length, and value (TLV) of DHCP options, and are converted to decimal to identify the DHCP options, as defined in RFC 2132.

The number of hex values that make up a particular DHCP option varies, depending on the length of the option. For example, the first DHCP option specified in the output includes three sets of hex values (35 01 01). The first hex value (35) identifies the option type, the second value (01) indicates the length of the value entry, which in this case is one set of hex values. The third hex value (01) specifies the value for the DHCP option.

In the second DHCP option specification (39 02 02 40), the hex value 39 is the type, and the length of 02 specifies that two sets of hex entries make up the value for the option. Therefore, this option specification uses four sets of hex entries; one for the type (39), one to specify the length (02), and two for the option value (02 40).

The third DHCP option is specified by the hex values 3d 07 01 00 10 94 00 00 08. The hex value 3d is the type, followed by the length (07), which specifies that the next seven sets of hex entries make up the value for the option. Therefore, this option specification uses a total of nine sets of hex entries; one for the type (3d), one to specify the length (07), and seven for the value of the DHCP option (01 00 10 94 00 00 08).

[Table 51 on page 1178](#) describes the first two options in more detail.

Table 51: DHCP Options Description

Option	Type	Length	Value
35 01 01	35 = decimal 53 (Code 53 in RFC 2132 is the DHCP Message Type option)	01 = the length of the option is one set of hex values (the next set in the list)	01 = value of the message type that is described in RFC 2132. The code 01 specifies a message type of DHCPDISCOVER.
39 02 02 40	39 = decimal 57 (Code 57 is the Maximum DHCP Message Size option)	02 = the length of the option is two sets of hex values (the next two sets in the list)	0240 = converted to a length of 576 octets

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, you can use the DHCPv6-Options VSA (26-207) to differentiate between different classes of subscribers during DHCPv6 relay authentication.
17.4R1	Starting in Junos OS Release 17.4R1, the exchange of parameters is enhanced with the introduction of several new vendor-specific attributes (VSAs) and changes to the existing DHCP-Options VSA (26-55).

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Standard and Vendor-Specific RADIUS Attributes | 3](#)

Common DHCP Configuration for Interface Groups and Server Groups

IN THIS SECTION

- [Grouping Interfaces with Common DHCP Configurations | 1179](#)
- [Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces | 1182](#)
- [Configuring Group-Specific DHCP Local Server Options | 1183](#)
- [Configuring Group-Specific DHCP Relay Options | 1184](#)
- [Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups | 1185](#)

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

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

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

2. Create the group and assign a name.

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

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the *interface interface-name* statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```

4. (Optional) You can use the *upto* option to specify a range of interfaces for a group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

5. (Optional) You can use the *exclude* option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
user@host# set interface fe-1/0/1.6 exclude
user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
```

Example- 2

To configure an interface group, use the *group* statement.

You can specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the `interface interface-name` statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. For example:

1. The extended DHCP applications enable you to group together a set of interfaces and apply a common DHCP configuration to the named interface group.

```
group boston {
    interface 192.168.10.1;
    interface 192.168.15.5;
}
```

2. You can use the `upto` option to specify a range of interfaces on which the extended DHCP application is enabled. For example:

```
group quebec {
    interface 192.168.10.1 upto 192.168.10.255;
}
```

3. You can use the `exclude` option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
group paris {
    interface 192.168.100.1 exclude;
    interface 192.168.100.100 upto 192.168.100.125 exclude;
}
```

Example:

```
group group-name {
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            mac-address;
            option-60;
```

```

        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix-string;
    }
}
interface interface-name <upto upto-interface-name> <exclude>;
}

```

Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following *configuration statement*:

```
user@host# set interface interface-name upto upto-interface-name
```

- The start subunit, interface *interface-name*, serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit .0 subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, interface *ge-2/2/2* is treated as interface *ge-2/2/2.0*.
- Ranged entries contain the *upto* option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a 0 (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because *ge-1/0/0.10* is the start subunit for both.

```

interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.10

```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface `ge-1/0/0.26` is common to all three.

```
dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
dhcp-local-server group ruby interface ge-1/0/0.26
dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.20 upto ge-1/0/0.40
```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

In the following example, the three ranges nest properly:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
```

- Discrete interfaces take precedence over ranges. In the following example, interface `ge-1/0/0.20` takes precedence and enforces an interface client limit of 5.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.15 upto ge-1/0/0.25 exclude
interface ge-1/0/0.20 overrides interface-client-limit 5
```

Configuring Group-Specific DHCP Local Server Options

You can include the following statements at the `[edit system services dhcp-local-server group group-name]` hierarchy level to set group-specific DHCP local server configuration options. Statements configured at the `[edit system services dhcp-local-server group group-name]` hierarchy level apply only to the named group of interfaces, and override any global DHCP local server settings configured with the same statements at the `[edit system services dhcp-local-server]` hierarchy level.

DHCPv6 local server supports the same set of statements with the exception of the `dynamic-profile` statement.

- `authentication` —Configure the parameters the router sends to the external AAA server.
- `dynamic-profile` —Specify the dynamic profile that is attached to a group of interfaces.
- `interface` —Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- `liveness-detection` —Configure bidirectional failure detection timers and authentication criteria for static routes, or Layer 2 liveness detection using ARP and Neighbor Discovery packets. For more information, see [DHCP Liveness Detection Overview](#).
- `overrides` —Override the default configuration settings for the extended DHCP local server. For information, see ["Overriding the Default DHCP Local Server Configuration Settings" on page 1034](#).
- `interface-tag`—(Optional) Specifies a tag name for the interface that will be associated with a DHCP configuration. Use the tag to identify subscribers associated with this DHCP local server group.

Configuring Group-Specific DHCP Relay Options

You can include the following statements at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level to set group-specific DHCP relay agent configuration options. Group-specific statements apply only to the named group of interfaces, and override any global DHCP relay agent settings for the same statement.

Include the statements at the `[edit forwarding-options dhcp-relay dhcpv6 group group-name]` hierarchy level to configure group-specific options for DHCPv6 relay agent.

- `active-server-group` —Configure an active server group to apply a common DHCP relay agent configuration for a named group of DHCP server addresses. For information, see ["Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups" on page 1185](#).
- `authentication` —Configure the parameters the router (or switch) sends to the external AAA server.
- `dynamic-profile` —Specify the dynamic profile that is attached to a group of interfaces.
- `interface` —Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- `liveness-detection` —Configure bidirectional failure detection timers and authentication criteria for static routes, or Layer 2 liveness detection using ARP and Neighbor Discovery packets. For more information, see [DHCP Liveness Detection Overview](#).

- `overrides` —Override the default configuration settings for the extended DHCP relay agent. For information, see ["Overriding the Default DHCP Relay Configuration Settings" on page 1036](#).
- `relay-agent-interface-id` —(DHCPv6 only) Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.
- `relay-agent-remote-id` —(DHCPv6 only) Insert the DHCPv6 Relay Agent Remote-ID option (option 37) in DHCPv6 packets destined for the DHCPv6 server.
- `relay-option` —Configure selective processing, which uses DHCP options in client packets to identify and filter client traffic, and to specify the action DHCP relay agent takes with the traffic. For more information, see ["Using DHCP Option Information to Selectively Process DHCP Client Traffic" on page 1054](#).
- `relay-option-82` —(DHCPv4 only) Enable or disable the insertion of option 82 information in packets destined for a DHCP server. For information, see ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).
- `service-profile` —Specify the default subscriber service, (or default profile) which is activated when the subscriber (or DHCP client) logs in and no other service is activated by a RADIUS server or a provisioning server. For more information, see ["Default Subscriber Service Overview" on page 1091](#).
- `interface-tag`—(Optional) Specifies a tag name for the interface that will be associated with a DHCP configuration. Use the tag to identify subscribers associated with this DHCP relay agent.

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups

You can apply a common DHCP or DHCPv6 relay configuration to a set of IP addresses configured as a server group. An active server group is sometimes referred to as a trusted group of servers.

You can configure active server groups globally or at the group level (configured with the `group`). When you apply the active server group at the group level, it overrides a global active server group configuration.

To configure a group of DHCP server addresses and apply them as an active server group:

1. Specify the name of the server group.

- For DHCPv4 servers:

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


- For DHCPv6 servers:

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

2. Add the IP addresses of the DHCP servers belonging to the group.

- For DHCPv4 servers:

```
[edit forwarding-options dhcp-relay server-group server-group-name]
user@host# set ip-address1
user@host# set ip-address2
```

- For DHCPv6 servers:

```
[edit forwarding-options dhcp-relay dhcpv6 active-server-group]
user@host# set ip-address1
user@host# set ip-address2
```



NOTE: Starting in Junos OS Release 18.4R1, up to 32 server IP addresses are supported per DHCPv4 server group. In earlier releases, a maximum of 5 server IP addresses are supported for DHCPv4 servers. Configuring more than the maximum number of server addresses results in a commit check failure.

3. Apply the server group as an active server group.

- At global level (DHCPv4)

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

- At group-level (DHCPv4)

```
[edit forwarding-options dhcp-relay group interface-group-name]
user@host# set active-server-group server-group-name
```

- At global level (DHCPv6)

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

- At group-level (DHCPv6)

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

Example: Configuring Active Server Groups in DHCP Relay Agent Configuration

For example, you might want to direct certain DHCP client traffic to a DHCP server. You can configure an interface group for each set of clients, specifying the DHCP relay interfaces for the group. In each of these groups, you specify an active server group to which each client groups traffic is forwarded. After a DHCP server group is created and server IP addresses are added to the group, the device used as the DHCP relay agent can forward messages to specific servers.

- Three groups of DHCP server addresses are configured, Default, Campus-A, and Campus-B.
- The Default group is applied as the global active server group for the overall DHCP relay configuration.
- The Campus-A server group is assigned as the active server group for interface group Campus-A-v10-DHCP-RELAY. DHCP traffic received on the interfaces in Campus-A-v10-DHCP-RELAY is forwarded to DHCP servers 198.51.100.100 and 198.51.100.101.
- The Campus-B server group is assigned as the active server group for interface group Campus-B-v200-DHCP-RELAY. DHCP traffic received on the interfaces in Campus-B-v200-DHCP-RELAY is forwarded to DHCP servers 198.51.100.55 and 198.51.100.56.
- All other DHCP traffic is forwarded to DHCP server 203.0.113.1.

```
[edit forwarding-options dhcp-relay]
#
# Server groups
user@host# set server-group Default 203.0.113.1
user@host# set server-group Campus-A 198.51.100.100
user@host# set server-group Campus-A 198.51.100.101
user@host# set server-group Campus-B 198.51.100.55
user@host# set server-group Campus-B 198.51.100.56
```

```
#
# Default server group applied globally.
user@host# set active-server-group Default
#
# Interface groups with application of active server groups
user@host# set group Campus-A-v10-DHCP-RELAY interface ge-1/1/0.1
user@host# set group Campus-A-v10-DHCP-RELAY interface ge-1/1/0.2
user@host# set group Campus-A-v10-DHCP-RELAY interface ge-1/1/0.3
user@host# set group Campus-A-v10-DHCP-RELAY active-server-group Campus-A
user@host# set group Campus-B-v200-DHCP-RELAY interface ge-2/2/0.4
user@host# set group Campus-B-v200-DHCP-RELAY interface ge-2/2/0.5
user@host# set group Campus-B-v200-DHCP-RELAY interface ge-2/1/0.6
user@host# set group Campus-B-v200-DHCP-RELAY active-server-group Campus-B
```

Note the following:

- In some configurations, servers in an active server group maintain redundant information about the DHCP clients. If the binding server later becomes inaccessible, the client is unable to renew the lease from that server. When the client attempts to rebind to a server, other servers in the group with the client information can reply with an ACK message. By default, instead of forwarding the ACK to the DHCP client, the relay agent drops any such ACKs that it receives from any server other than the binding server because the new server address does not match the expected server address in the DHCP client entry. Consequently the lease cannot be extended by any of the redundant servers.
- Starting in Junos OS Release 16.2R1, you can enable a DHCPv4 relay agent to forward DHCP request (renew or rebind) ACKs from any server in the active server group (thus, any trusted server). The relay agent updates the client entry with the new server address. Because the servers in the group are expected to mirror the client information exactly, the lease option is expected to be the same as for the original server and the relay agent does not need to verify the lease option.
- Starting in Junos OS Release 18.4R1, this capability is extended to allow a DHCP relay agent to forward DHCP information request (DHCPINFORM) ACK messages from any server in the active server group.

To enable ACK forwarding from any server in the active server group:

- Enable forwarding for the active server group.

```
[edit forwarding-options dhcp-relay active-server-group]
user@host# set allow-server-change
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1
16.2R1	Starting in Junos OS Release 16.2R1

RELATED DOCUMENTATION

[DHCP Overview](#) | [1015](#)

[DHCPv6 Local Server](#) | [1449](#)

[DHCPv6 Relay Agent](#) | [1455](#)

Number of DHCP Clients Per Interface

IN THIS SECTION

- [Specifying the Maximum Number of DHCP Clients Per Interface](#) | [1189](#)
- [Allowing Only One DHCP Client Per Interface](#) | [1191](#)

Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.



NOTE: The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the `interface-client-limit` statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.

- For DHCP local server:

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

- For DHCPv6 local server:

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

- 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
```

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the `interface-client-limit` statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```



NOTE: For DHCP local server and DHCP relay agent, you can use either the `interface-client-limit` statement or the `client-discover-match incoming-interface` statement to set a limit of one client per interface. The `interface-client-limit` statement with a value of 1

retains the existing client and rejects any new client connections. The `client-discover-match incoming-interface` statement deletes the existing client and allows a new client to connect.

Allowing Only One DHCP Client Per Interface

Subscriber management provides two methods that you can use to configure DHCP local server and DHCP relay agent to allow only one DHCP client per interface. The two methods differ on which client is allowed on the interface—the new client or the existing client. The two methods are supported by both DHCP local server and DHCP relay agent, and can be configured globally, for a group of interfaces, or for a specific interface.

- Accept new client—Delete the existing client binding and allow the new client to connect. To configure this action, use the ... overrides `client-discover-match incoming-interface` statement.
- Keep existing client—Retain the existing client binding on the interface and reject any requests from new DHCP clients. To configure this action, use the ... overrides `interface-client-limit 1` statement to specify a maximum of one client.

To configure the router to delete the existing client binding on the interface and allow the new client to connect:

1. Specify that you want to configure override options.

- For DHCP local server:

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

- For DHCP relay agent:

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

2. Configure the router to view all client connections on the interface as coming from the same client, which allows a new client to replace the existing client. This example shows the DHCP local server configuration. The DHCP relay agent configuration is similar.

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match incoming-interface
```

To configure the router to keep the existing client binding on the interface and refuse connections from new clients:

1. Specify that you want to configure override options.

- For DHCP local server:

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

- For DHCP relay agent:

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

2. Set the maximum number of clients allowed per interface to one. This example shows the DHCP local server configuration. The DHCP relay agent configuration is similar.

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit 1
```

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[Conserving IP Addresses Using DHCP Auto Logout | 1208](#)

Maintaining DHCP Subscribers During Interface Delete Events

IN THIS SECTION

- [Maintaining Subscribers During Interface Delete Events | 1193](#)

- [Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | 1194](#)
- [Configuring an ACX Series DHCP Local Server to Preserve Subscriber Binding Information | 1194](#)
- [Verifying and Managing DHCP Subscriber Binding During Interface Delete Events | 1196](#)

Maintaining Subscribers During Interface Delete Events

IN THIS SECTION

- [Benefits of Maintaining Subscriber Bindings | 1194](#)

You can configure the router to maintain DHCP subscribers (maintain the subscriber bindings) when an event occurs that normally results in the router deleting the subscriber. For example, by default, the router logs out DHCP subscribers when an interface delete event occurs, such as a line card reboot or failure. However, if you configure the router to maintain subscribers, the router identifies each subscriber that was on the deleted interface, and resumes normal packet processing for the subscriber when the interface is restored. This procedure does not maintain subscribers that are deleted during router reboots or failures.



NOTE: Subscribers are logged off as usual when their lease expires, even if the router is configured to maintain subscribers and the subscriber is on a deleted interface that has not yet been restored.

You configure the router to maintain subscribers on a global basis—the configuration applies to DHCP local server, DHCPv6 local server, and DHCP relay clients in all logical routers and routing instances. When you enable the maintain subscribers feature, the router applies the feature to existing subscribers as well as subscribers who later connect.

If the maintain subscribers feature is enabled on the router, you can explicitly delete a subscriber binding and log out the subscriber by either specifying a lease expiration timeout or using one of the following commands, as appropriate:

- `clear dhcp server binding`
- `clear dhcpv6 server binding`
- `clear dhcp relay binding`

Benefits of Maintaining Subscriber Bindings

Reduces the time to restore the subscriber session and minimizes loss of subscriber service.

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events

You can specify a configuration in which the router does not log out a subscriber when the subscriber's interface is deleted.



NOTE: This procedure does not maintain subscribers during router reboots or failures.

To configure the router to maintain DHCP subscribers when the subscriber interface is deleted:

1. Specify that you want to configure subscriber management.

```
[edit system services]
user@host# edit subscriber-management
```

2. Configure the router to support the maintain-subscriber feature.

```
[edit system services subscriber-management]
user@host# edit maintain-subscriber
```

3. Configure the router to enable the maintain-subscriber feature when an interface-delete event occurs.

```
[edit system services subscriber-management maintain-subscriber]
user@host# set interface-delete
```

Configuring an ACX Series DHCP Local Server to Preserve Subscriber Binding Information

When an ACX series router functioning as a DHCP local server reboots, by default, all the subscriber binding information is lost. You can configure the local server to preserve the subscriber binding information to a file in **/var/preserve**. When the router reboots, the DHCP local server reads the file and restores the subscriber binding information and resumes normal packet processing for the subscriber. By default, a new file is generated every 24 hours from the commit time, but you can specify a backup interval of 1 through 48 hours. The configuration is a global setting for each routing instance.

To configure an ACX Series DHCP local server to store subscriber binding information:

- Enable persistent storage.

```
[edit system services dhcp-local-server]
user@host# set persistent-storage automatic
```

To configure a file to store subscriber binding information:

1. Specify a filename for storing subscriber binding information. By default, the file is named `jdhcpd_client_data`.



NOTE: A commit error occurs if you try to configure the file with a name that is already present in the `/var/preserve` directory.

```
[edit system processes dhcp-service]
user@host# set persistent-storage file-name
```

2. Specify a frequency to back up the file.

```
[edit system processes dhcp-service]
user@host# set persistent-storage backup-interval hours
```

The following example saves the binding information to `/var/preserve/acx-local-server1-client-data` every 8 hours:

```
[edit system processes dhcp-service]
user@host# set persistent-storage acx-local-server1-client-data backup-interval 8
```

SEE ALSO

[Extended DHCP Local Server Overview](#)

[DHCP Local Server Handling of Client Information Request Messages](#)

[DHCP Duplicate Client Differentiation Using Client Subinterface Overview](#)

[Guidelines for Configuring Support for DHCP Duplicate Clients](#)

[Configuring DHCP Client-Specific Attributes](#)

[Automatically Logging Out DHCP Clients](#)

[Enabling Processing of Client Information Requests](#)[Configuring a DHCP Client on ACX Series](#)

Verifying and Managing DHCP Subscriber Binding During Interface Delete Events

IN THIS SECTION

- [Purpose | 1196](#)
- [Action | 1196](#)

Purpose

Display information related to the DHCP maintain-subscribers feature and explicitly log out maintained clients.

Action

- To display DHCP local server binding information for the DHCP maintain subscribers feature:

```
user@host>show dhcp server binding detail
```

- To display DHCPv6 local server binding information for the DHCP maintain subscribers feature:

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

- To display DHCP relay binding information for the DHCP maintain subscribers feature:

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

- To explicitly log out a DHCP local server subscriber when the maintain subscriber feature is enabled:

```
user@host>clear dhcp server binding binding-type
```

- To explicitly log out a DHCPv6 local server subscriber when the maintain subscriber feature is enabled:

```
user@host>clear dhcpv6 server binding binding-type
```

- To explicitly log out a DHCP relay subscriber when the maintain subscriber feature is enabled:

```
user@host>clear dhcp relay binding binding-type
```

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[Conserving IP Addresses Using DHCP Auto Logout | 1208](#)

Dynamic Reconfiguration of Clients From a DHCP Local Server

IN THIS SECTION

- [Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | 1198](#)
- [Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 1201](#)
- [Configuring Dynamic Reconfiguration Attempts for DHCP Clients | 1202](#)
- [Configuring Deletion of the Client When Dynamic Reconfiguration Fails | 1203](#)
- [Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | 1204](#)
- [Configuring a Token for DHCP Local Server Authentication | 1204](#)

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients

IN THIS SECTION

- [Default Client/Server Interaction | 1198](#)
- [Dynamic Client/Server Interaction for DHCPv4 | 1199](#)
- [Dynamic Client/Server Interaction for DHCPv6 | 1199](#)
- [Manually Forcing the Local Server to Initiate the Reconfiguration Process | 1200](#)
- [Action Taken for Events That Occur During a Reconfiguration | 1200](#)
- [Benefits of Dynamic Reconfiguration of DHCP Local Server Clients | 1201](#)

Dynamic reconfiguration of clients enables the extended DHCP local server to initiate a client update without waiting for the client to initiate a request.

Default Client/Server Interaction

Typically the DHCP client initiates all of the basic DHCP client/server interactions. The DHCP server sends information to a client only in response to a request from that client. This behavior does not enable a client to be quickly updated with its network address and configuration in the event of server changes:



NOTE: Technically, the DHCP client/server interactions are the same on routers and switches. However, the primary usage of this technology on the routers is for subscriber management. The switches are not used for subscriber management. Therefore, this topic provides two sample scenarios. The actions are the same, but the implementation details are different.

- On routers—Suppose a service provider restructures its addressing scheme or changes the server IP addresses that it provided to clients. Without dynamic reconfiguration, the service provider typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, the provider can wait for customers to make a service call about the network failures and then instruct them to power cycle their customer premises equipment to reinitiate the connection. Neither of these actions is timely or convenient for customers.

- On switches—Suppose you restructure the addressing scheme or change the server IP addresses that the DHCP server provides to clients. Without dynamic reconfiguration, the network typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, you can wait for users to notify you of the network failures and then instruct them to power cycle their equipment to reinitiate the connection. Neither of these actions is timely or convenient for users.

Dynamic Client/Server Interaction for DHCPv4

Dynamic reconfiguration for DHCPv4 is available through a partial implementation of RFC 3203, *DHCP Reconfigure Extension* for DHCPv4. It enables the DHCPv4 local server to send a message to the client to force reconfiguration.

The server sends a *forcerenew* message to a DHCPv4 client, initiating a message exchange. In response, DHCPv4 clients that support the *forcerenew* message then send a lease renewal message to the server. The server rejects the lease renewal request and sends a NAK to the client, causing the client to reinitiate the DHCP connection. A successful reconnection results in the reconfiguration of the DHCP client. Only the exchange of *forcerenew*, *renew*, and NAK messages is supported from RFC 3202. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to *forcerenew* messages other than to forward them to the client.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a *forcerenew* message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber services or DHCP-managed services, such as forwarding and statistics, continue to work. Client statistics are not maintained in the interval between a successful reconfiguration and the subsequent client binding. When the server responds to the client renewal request with a NAK, the client entry is removed from the binding table and final statistics are reported. New statistics are collected when the client sends a discover message to establish a new session.

Dynamic Client/Server Interaction for DHCPv6

Dynamic reconfiguration for DHCPv6 is available through a partial implementation of RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. It enables the DHCPv6 local server to send a message to the client to force reconfiguration.

DHCPv6 servers send *reconfigure* messages to DHCPv6 clients, initiating a message exchange. In response, DHCPv6 clients that support the *reconfigure* message transition to the renewing state and send a *renew* message to the server. The server returns a reply message with a lifetime of zero (0). The client transitions to the init state and sends a *solicit* message. The server sends an *advertise* message to

indicate that it is available for service. The client sends a request for configuration parameters, which the server then includes in its reply. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to reconfigure messages other than to forward them to the client.

When a DHCPv6 server is triggered to initiate reconfiguration on a bound DHCPv6 client, the client transitions to the reconfigure state. All subscriber services, such as forwarding and statistics, continue to work. The server then sends the reconfigure message to the client. If the DHCPv6 client is already in the reconfigure state, the DHCPv6 server ignores the reconfiguration trigger. For clients in any state other than bound or reconfigure, the server clears the binding state of the client, as if the `clear dhcpv6 server binding` command had been issued.

Manually Forcing the Local Server to Initiate the Reconfiguration Process

You can force the local server to initiate the reconfiguration process for clients by issuing the `request dhcp server reconfigure` command for DHCPv4 clients, and the `request dhcpv6 server reconfigure` command for DHCPv6 clients. Command options determine whether reconfiguration is then attempted for all clients or specified clients.

Action Taken for Events That Occur During a Reconfiguration

Events that take place while a reconfiguration is in process take precedence over the reconfiguration. [Table 52 on page 1200](#) lists the actions taken in response to several different events.

Table 52: Action Taken for Events That Occur During a Reconfiguration

Event	Action
Server receives a discover (DHCPv4) or solicit (DHCPv6) message from the client.	Server drops packet and deletes client.
Server receives a request, renew, rebind, or init-reboot message from the client.	DHCPv4—Server sends NAK message and deletes client. DHCPv6—Server drops packet and deletes client. Server replies to renew message with lease time of zero (0).
Server receives a release or decline message from the client.	Server deletes client.
The client lease times out.	Server deletes client.

Table 52: Action Taken for Events That Occur During a Reconfiguration (*Continued*)

Event	Action
The <code>clear dhcp server binding</code> command is issued.	Server deletes client.
The <code>request dhcp server reconfigure (DHCPv4)</code> or <code>request dhcpv6 server reconfigure (DHCPv6)</code> command is issued.	Command is ignored.
GRES or DHCP restart occurs.	Reconfiguration process is halted.

Benefits of Dynamic Reconfiguration of DHCP Local Server Clients

- Enable the DHCP local server to dynamically reconfigure DHCP clients, avoiding extended outages because of server configuration changes that otherwise require the server to wait for the client to renew its lease or rebind to the server.

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

1. Enable dynamic reconfiguration with default values for all clients.

For DHCPv4:

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

For DHCPv6:

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

2. (Optional) Enable dynamic reconfiguration for only the clients serviced by a group of interfaces.

For DHCPv4:

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


For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 group group-name]
user@host# set reconfigure
```

3. (Optional) Configure how the server attempts reconfiguration.
See ["Configuring Dynamic Reconfiguration Attempts for DHCP Clients" on page 1202](#).
4. (Optional) Configure the response to a failed reconfiguration.
See ["Configuring Deletion of the Client When Dynamic Reconfiguration Fails" on page 1203](#).
5. (Optional) Configure the behavior in response to a RADIUS-initiated disconnect.
See ["Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect" on page 1204](#).
6. (Optional) Configure a token for rudimentary server authentication.
See ["Configuring a Token for DHCP Local Server Authentication" on page 1204](#).
7. (Optional) Prevent DHCPv6 clients from binding if they do not support reconfigure messages.
See ["Preventing Binding of Clients That Do Not Support Reconfigure Messages" on page 1451](#).

Configuring Dynamic Reconfiguration Attempts for DHCP Clients

You can configure how many attempts the local server makes to initiate reconfiguration of the DHCP client by sending `forcerenew` or `reconfigure` messages. You can also specify how long the server waits between attempts. By default, eight attempts are made and the initial interval is two seconds.

Each successive attempt doubles the interval between attempts. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on. A group configuration takes precedence over a DHCP local server configuration.

(Optional) To configure DHCP local server reconfiguration behavior for all DHCP clients:

1. Specify the number of reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set attempts 5
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set attempts 5
```

2. Specify the interval between reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set timeout 8
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set timeout 8
```

To override the global configuration for a particular group of clients, include the statements at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

Configuring Deletion of the Client When Dynamic Reconfiguration Fails

You can configure the local server to delete the client when the maximum number of reconfiguration attempts has been made without success. By default, the client's original configuration is restored.

(Optional) To configure the DHCP local server to delete the client when reconfiguration is not successful, for all clients:

- Specify the client deletion.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set clear-on-terminate
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set clear-on-terminate
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect

You can configure the local server to reconfigure the client when the client receives a RADIUS-initiated disconnect. By default, the client is deleted when a RADIUS-initiated disconnect is received.

(Optional) To configure the DHCP local server to reconfigure the client instead of deleting the client when a RADIUS-initiated disconnect is received, for all clients:

- Specify the RADIUS-initiated disconnect trigger.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure trigger]
user@host# set radius-disconnect
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure trigger]
user@host# set radius-disconnect
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure trigger]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure trigger]` hierarchy level.

Configuring a Token for DHCP Local Server Authentication

You can configure an authentication token to provide rudimentary protection against inadvertently instantiated DHCP servers. You can configure the local server to include a constant, unencoded token in the DHCP forcere Renew message as part of the authentication option it sends to clients. If the service provider has previously configured the DHCP client with a token, then the client can compare that token against the newly received token. If the tokens do not match, the DHCP client discards the forcere Renew message. This functionality corresponds to RFC 3118, *Authentication for DHCP Messages*, section 4.

(Optional) To configure the DHCP local server to include a token in the forcere Renew message sent to the client, for all clients:

- Specify the token.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set token token-value
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set token token-value
```

(Optional) For only a particular group of clients:

- Specify the token.

For DHCPv4:

```
[edit system services dhcp-local-server group group-name reconfigure]
user@host# set token token-value
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
user@host# set token token-value
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.1	Starting in Junos OS Release 24.4R1, the existing JUNOS DHCPV4/V6 Server 'FORCENEW/RECONFIGURE' support now also supports a 'deferred-NAK' option, whereby if the the DHCP client does not immediately respond to the FORCE-RENEW/RECONFIGURE request (or any of its subsequent limited retries), the subscriber session is left in place with full connectivity and the session gets flagged for 'deferred-NAK'. This session state is mentioned persistently across daemon restarts and GRES/ISSU events.

RELATED DOCUMENTATION

DHCP Overview	1015
DHCPv6 Local Server	1449
DHCPv6 Relay Agent	1455
DHCP Monitoring and Management	1224

Understanding Deferred NACK on DHCP Reconfigure Abort

IN THIS SECTION

- [Benefits of Deferred NACK on DHCP Reconfigure Abort | 1206](#)
- [Overview | 1207](#)
- [Configuration Commands | 1207](#)

The Deferred NACK on DHCP Reconfigure Abort feature enhances DHCPv4/DHCPv6 Local Server Reconfigure functionality by maintaining DHCP bindings when a client does not support or respond to Force-Renew/Reconfigure requests. By using the `defer-nack-on-abort` configuration option, you can mark a binding for deferred NACK, ensuring uninterrupted client connectivity while scheduling a NACK response for the next RENEW cycle. This feature is configured through specific CLI commands within the `dhcp-local-server` hierarchy and is mutually exclusive with the `clear-on-abort` option. It integrates seamlessly with existing operational and show commands to manage and monitor DHCP bindings effectively, providing more graceful handling of reconfiguration failures and minimizing service disruptions.

Benefits of Deferred NACK on DHCP Reconfigure Abort

- Minimizes service disruptions by maintaining client connectivity even when a DHCP client does not support or respond to reconfigure requests.
- Allows for graceful handling of reconfiguration failures by deferring the NACK response to the next RENEW cycle, ensuring the client can renegotiate and reestablish its session without immediate disconnection.
- Provides network operators with more control and flexibility in managing DHCP bindings, reducing the need for manual interventions and improving operational efficiency.
- Ensures compatibility with existing DHCP server functionalities and infrastructure, maintaining consistent behavior across the network.
- Enhances the reliability and stability of the DHCP service by supporting clients that do not implement the reconfigure feature, leading to a more robust and fault-tolerant network environment.

Overview

The Deferred NACK on DHCP Reconfigure Abort feature extends the DHCPv4/DHCPv6 Local Server Reconfigure functionality by ensuring that DHCP bindings are maintained even when a client does not respond to Force-Renew/Reconfigure requests. By configuring the `defer-nack-on-abort` option, you can defer the Negative Acknowledgement (NACK) until the client initiates its next RENEW cycle. This enables clients to maintain their connectivity and only forces a session renegotiation when the client is ready, minimizing service disruptions.

To implement this feature, use specific CLI commands within the `dhcp-local-server` hierarchy. The primary command to enable this functionality is `set system services dhcp-local-server reconfigure defer-nack-on-abort`. This option is mutually exclusive with the `clear-on-abort` option, which would otherwise immediately clear the binding if the client fails to respond. By using the `defer-nack-on-abort` configuration, the DHCP server marks the binding to send a NACK during the next client-initiated RENEW cycle, ensuring a smooth transition without abrupt disconnections.

Verify the deferred NACK state using existing show commands such as `show dhcp server binding detail` and `show dhcpv6 server binding detail`. These commands provide detailed information about DHCP bindings and indicate if the `defer-nack-on-abort` status is active. This feature is supported across FreeBSD and Linux platforms, ensuring backward compatibility and compliance with current security practices. It also integrates seamlessly with the existing DHCP infrastructure, maintaining consistent behavior and operational stability.

Configuration Commands

To configure the Deferred NACK on DHCP Reconfigure Abort feature, use the following CLI commands:

For DHCPv4 bindings:

```
set system services dhcp-local-server reconfigure defer-nack-on-abort set system services dhcp-local-server group
<grp-name> reconfigure defer-nack-on-abort
```

For DHCPv6 bindings:

```
set system services dhcp-local-server dhcpv6 reconfigure defer-nack-on-abort set system services dhcp-local-server
dhcpv6 group <grp-name> reconfigure defer-nack-on-abort
```

These commands are placed within the `dhcp-local-server` hierarchy and allow you to configure the `defer-nack-on-abort` option for specific groups or the entire DHCP service. This flexibility enables you to tailor the DHCP server behavior to meet specific network requirements while ensuring minimal service disruption and optimized client handling.

Conserving IP Addresses Using DHCP Auto Logout

IN THIS SECTION

- [DHCP Auto Logout Overview | 1208](#)
- [Automatically Logging Out DHCP Clients | 1210](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout | 1211](#)
- [DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers | 1212](#)
- [Automatically Logging Out DHCPv6 Clients | 1213](#)

DHCP Auto Logout Overview

IN THIS SECTION

- [Auto Logout Overview | 1208](#)
- [How DHCP Identifies and Releases Clients | 1209](#)
- [Option 60 and Option 82 Requirements | 1210](#)

This topic provides an introduction to the DHCP auto logout feature and includes the following sections:

Auto Logout Overview

Auto logout is supported for DHCP local server and DHCP relay agent. It improves the efficiency of DHCP IP address assignment by allowing IP addresses to be immediately released and returned to the address pool when DHCP clients are no longer using the addresses. DHCP can then assign the addresses to other clients. Without auto logout, an IP address is blocked for the entire lease period, and DHCP must wait until the address lease time expires before reusing the address.

Auto logout is particularly useful when DHCP uses long lease times for IP address assignments and to help avoid allocating duplicate IP addresses for a single client.

For example, you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address— the previous IP address assigned to the client (the old STB) remains

blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

How DHCP Identifies and Releases Clients

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier, and subnet. However, in some cases this type of identification might not be sufficient. For example, in the previous STB example, each STB has a different MAC address, so DHCP incorrectly assumes that an upgraded or replacement STB is a new client.

In order to explicitly identify clients, auto logout uses a secondary identification method when the primary identification method is unsuccessful— the primary method is considered unsuccessful if the MAC address or Client Identifier does not match that of an existing client. Subscriber management supports two secondary identification methods that you can configure.

- Incoming interface method— DHCP views a new client connection on the interface as if it comes from the same client. DHCP deletes the existing client binding before creating a binding for the newly connected device. This method allows only one client device to connect on the interface.



NOTE: The incoming interface method differs from the `overrides interface-client-limit 1` statement, which retains the existing binding and rejects the newly connected client.

- Option 60 and option 82 method— DHCP considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the Discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

If the DHCP relay receives a Discover message from an existing client, the DHCP relay forwards the Discover message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.



NOTE: If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.

Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, the DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in [DHCP Relay Agent Option 82 Value for Auto Logout](#).

Automatically Logging Out DHCP Clients

You can configure the extended DHCP local server and extended DHCP relay to automatically log out DHCP clients. Auto logout immediately releases an existing client when DHCP receives a discover packet from a client whose identity matches an existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.



NOTE: When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure DHCP client auto logout:

1. Specify that you want to configure override options.

- For DHCP local server:

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

- For DHCP relay agent:

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

2. Enable auto logout and specify the secondary identification method you want to use when the primary identification method is unsuccessful.

- For example, to configure DHCP local server to use the incoming interface method:

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match incoming-interface
```

- For example, to configure DHCP relay agent to use the option 60 and option 82 method:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set client-discover-match option60-and-option82
```



NOTE: If you change the auto logout configuration, existing clients continue to use the auto logout setting that was configured when they logged in. New clients use the new setting.

How DHCP Relay Agent Uses Option 82 for Auto Logout

Table 53 on page 1211 indicates how the DHCP relay agent determines the option 82 value used for the client auto logout feature. Depending on the configuration settings, DHCP relay agent takes the action indicated in the Action Taken column.

Table 53: DHCP Relay Agent Option 82 Value for Auto Logout

DHCP Relay Agent Configuration Settings				giaddr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option- 82"	Override "always-write-option-82"		
No	No	–	–	–	No secondary search performed
No	Yes	Yes	–	–	Use option 82 from packet

Table 53: DHCP Relay Agent Option 82 Value for Auto Logout (*Continued*)

DHCP Relay Agent Configuration Settings				giaddr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override “trust-option- 82”	Override “always-write-option-82”		
No	Yes	No	–	Zero	Drop packet
No	Yes	No	–	Non-zero	Use option 82 from packet
Yes	No	–	–	–	Use configured option 82
Yes	Yes	No	–	Zero	Drop packet
Yes	Yes	No	No	Non-zero	Use option 82 from packet
Yes	Yes	No	Yes	Non-zero	Overwrite the configured option 82
Yes	Yes	Yes	No	–	Use option 82 from packet
Yes	Yes	Yes	Yes	–	Overwrite the configured option 82

DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers

By default, the DHCPv6 local server and the DHCPv6 relay agent identify clients on the basis of the client identifier. The DHCPv6 local server and the DHCPv6 relay agent can also identify a DHCPv6 client by the incoming interface. You use the `incoming-interface` option with the `client-negotiation-match` statement so that only one client device connects on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

For example, consider an environment that includes a set-top box (STB) or any other such customer premises equipment (CPE) device configured to get configuration information from the DHCPv6 server. In the network configuration, one CPE device is supported over an interface. The DHCPv6 server is configured to provide the CPE devices with long lease timers. If the CPE device is disconnected for repair or upgraded, the new CPE device goes through the DHCPv6 Solicit process to receive the configuration information from the DHCPv6 server. Because the client identifier is different from that of the previous device, the DHCPv6 local server or the DHCPv6 relay agent treats the DHCPv6 Solicit message as a new client and adds the new binding. Because the old device might not gracefully log out, the old IP address is not released until the lease expires.

If the `client-negotiation-match incoming-interface` statement is configured, on receiving a DHCPv6 Solicit message, the DHCPv6 clients are searched on the basis of their client identifiers and the incoming interface option. If an existing DHCPv6 client binding is found based on the match criteria, the binding is removed and the new client is processed. If the old CPE device is disconnected and a DHCPv6 Solicit message is received for the new CPE device, the feature uses the incoming interface to identify the client and remove the binding of the old CPE device, which allows for the release of the old IP address. The binding of the new CPE device replaces the old binding.

Automatically Logging Out DHCPv6 Clients

You can configure the extended DHCPv6 local server and extended DHCPv6 relay agent to automatically log out DHCPv6 clients based on DHCPv6 subscriber-match criteria. The automatic logout feature immediately releases an existing client when DHCPv6 receives a Solicit packet from a client whose incoming interface matches that of an existing client. DHCPv6 then releases the existing client IP address without waiting for the normal lease expiration.



NOTE: When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure automatic logout of DHCPv6 clients:

1. Specify that you want to configure override options to override the default configuration settings.
 - For the DHCPv6 local server:

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

- For the DHCPv6 relay agent:

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

2. Enable automatic logout and specify the incoming interface as the secondary identification method you want to use.

- For the DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set client-negotiation-match incoming-interface
```

- For the DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set client-negotiation-match incoming-interface
```



NOTE: If you change the automatic logout configuration, existing clients continue to use the automatic logout setting that was configured when they logged in. New clients use the new setting.

RELATED DOCUMENTATION

[Number of DHCP Clients Per Interface | 1189](#)

[DHCP Monitoring and Management | 1224](#)

[Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 1033](#)

[DHCP Overview | 1015](#)

[DHCPv6 Local Server | 1449](#)

[DHCPv6 Relay Agent | 1455](#)

DHCP Short Cycle Protection

IN THIS SECTION

- [DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions | 1215](#)
- [Configuring DHCP Short-Cycle Protection | 1218](#)
- [Verifying and Managing DHCP Short-Cycle Protection | 1222](#)

DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions

IN THIS SECTION

- [Conditions That Can Cause Failed or Short-Lived DHCP Client Sessions | 1216](#)
- [How DHCP Short-Cycle Protection Works | 1216](#)
- [Termination of the Lockout Condition | 1217](#)
- [Benefits of Using DHCP Short Cycle Protection | 1218](#)

In highly scaled networks, a significant number of DHCP client negotiations fail before the session is established, resulting in high loading on the router and external authentication servers. Some CPE devices automatically retry negotiation on failure, some with very short retry intervals. A malicious client might mount an authentication attack by sending repeated, frequent login requests. These events can result in a significant load on the router and the external authentication server.

Starting in Junos OS Release 18.2R1, *DHCP short cycle protection*, also called *DHCP client lockout*, enables the router to reduce these loads by identifying and temporarily locking out clients that continually fail negotiation and have short negotiation cycles as well as clients that frequently complete connections but log out soon after logging in.

Identified clients are prevented from access by temporarily locking them out for an exponentially increasing *lockout period*. The router drops DHCP discover or solicit messages from these clients while they are locked-out. The router tracks clients by the client identifier for DHCPv4 clients or DHCP unique identifier (DUID) for DHCPv6 clients. Both types of client identifiers can be referred to as client keys. The client key enables the DHCP server to associate a client with its lease and configuration parameters. Using the client key for DHCP short-cycle protection tracking enables the router to prevent

one client from negotiating a session while allowing other clients using the same logical interface to successfully negotiate sessions.

The initial lockout period for a client has a short duration. The goal here is to not negatively affect legitimate clients, for example, those that fail just once or that log in periodically to check their email and then log out again. By targeting clients that continually fail negotiation or log in and out frequently at short intervals, short-cycle protection reduces both the connection processing load on the router and the authentication load on external authentication servers. It has the effect of improving throughput by deferring client sessions that do not make progress in favor of sessions that complete.

Conditions That Can Cause Failed or Short-Lived DHCP Client Sessions

Conditions that can cause a failed or short-lived client session include:

- Authentication denials from external AAA servers, such as RADIUS or Diameter, due to the absence of a corresponding entry in the RADIUS database or due to improper login attempts.
- Router or external authentication server unreachability due to network failure or misconfiguration.
- Insufficient memory resources to create a dynamic subscriber interface.
- Protocol negotiation failures with the CPE.
- Client logout shortly after a successful login; this action creates a fully negotiated and configured client session before the session is torn down.

How DHCP Short-Cycle Protection Works

DHCP short-cycle protection is disabled on the router by default. When you enable it by including the `short-cycle-protection` statement at a global, group, or interface level, the router does the following for DHCP sessions on static and dynamic logical interfaces:

1. Detects short-lived client sessions, also referred to as *short-cycle events*, and locks out the client based on the following events:
 - E0: Time when `jdhcpd` declares the client session to be active.
 - E1: Time when `jdhcpd` declares the client session should be torn down.
 - E2: Time when `jdhcpd` deletes the client session entry from the database.

A short-cycle event occurs when the interval between E0 and E1 is less than or equal to 60 seconds. When the interval is greater than 60 seconds, the logout is considered normal. If the router declares the session to be short-lived, it adds the client to the lockout database at time E2.

2. Temporarily locks out the specified DHCP client by preventing connection to the router.

During lockout, the router drops negotiation packets (DHCP discover and solicit messages) from the client until the lockout period expires. When the lockout period expires, the client can resume normal negotiation of the connection.

You can set a range for the lockout period by specifying a minimum and maximum length with the `short-cycle-protection` statement. You must specify both a minimum and a maximum value.

3. Tracks the time between a client's repeated short-cycle events to determine whether to increase the lockout time for a subsequent short-cycle event. The interval between events is compared to the *grace time threshold*. By default, the grace time threshold is 900 seconds, but it is automatically set to the maximum lockout time if that value is greater than 900 seconds.

If no subsequent negotiation is attempted within the grace time, the client entry is removed from the lockout database.

If a subsequent negotiation is attempted before the grace threshold is reached, it is treated as another short-cycle event and the lockout penalty is increased. The penalty is increased exponentially each time the negotiation is attempted within the grace time.

The initial lockout period is based on the configured minimum value. Additional penalties are calculated as follows, where n is the number of consecutive short-cycle events that occur within the grace time:

$$\text{Lockout time} = (\text{Lockout minimum time}) \times [2^{(n-1)}]$$

For example, with a minimum duration of 1 second and a maximum duration of 300 seconds, the initial lockout period is 1 second; subsequent penalties increase to 2 seconds, then 4 seconds, 8 seconds, 16 seconds, 32 seconds, 64 seconds, 128 seconds, 256 seconds and finally 300 seconds. The final lockout period is 300 seconds instead of 512 seconds because no penalty can exceed the maximum value of the lockout range.

If the lockout time reaches the maximum, then it stays at that value for each subsequent lockout period until the time between short-cycle events is greater than the grace threshold.

Termination of the Lockout Condition

When a DHCP client is locked out, the lockout condition persists until all lockout timers have expired, *except* when any of the following occurs:

- You administratively clear the lockout condition by issuing one of the following operational commands:
 - `clear dhcp relay lockout-entries`
 - `clear dhcp server lockout-entries`
 - `clear dhcpv6 relay lockout-entries`

- `clear dhcpv6 server lockout-entries`
- You reset the FPC on which the client session undergoing lockout is configured.
- You reset the Routing Engine.

When any of these events occurs, `jdhcpd` terminates lockout and clears the lockout history for all affected client sessions. The released clients are allowed to negotiate again. Because there is no retained history, the lockout period starts with the minimum value if a subsequent short-cycle event occurs for one of these clients.

When a dynamic VLAN or demux VLAN logical interface is removed from an underlying physical interface that is configured with `remove-when-no-subscribers`, the lockout of affected clients persists until all the timers have expired. If the logical interface is recreated before all timers expire, then the lockout state is applied to the re-created logical interfaces.

Benefits of Using DHCP Short Cycle Protection

- Reduces excessive control plane loading on the router and authentication, authorization, and provisioning loading on the external authority server.
- Reduces the resources required to process DHCP control packets and to negotiate and terminate short-lived connections.
- Temporarily defers subsequent attempts for clients with failed or short-lived client sessions in favor of sessions can complete successfully and last for more than a short duration.
- Reduces the resources required to authenticate and terminate these connections on external authentication servers, such as RADIUS and Diameter.
- Enables lockout of a single failed or short-lived DHCP session without disrupting other DHCP sessions on the same interface.

Because DHCP short-cycle protection identifies each client session by its unique client ID, the router can lock out only the offending DHCP client while enabling other DHCP clients on the same interface to successfully negotiate the connection.

Configuring DHCP Short-Cycle Protection

In highly scaled networks, a significant number of DHCP client negotiations fail before the session is established, resulting in high loading on the router and external authentication servers. You can enable DHCP short cycle protection on the router to identify DHCP clients that either login frequently and briefly or continually fail to connect, then lock the clients out from access and drop subsequent requests from these clients until a lockout timer expires. For clients that repeatedly log in frequently and briefly, the initial lockout time is short enough to have no noticeable impact. As these brief logins continue, the lockout period is exponentially increased. By targeting clients that continually fail negotiation or log in

and out frequently at short intervals, short-cycle protection reduces the connection processing load on the router and the authentication, authorization, and provisioning load on external authentication servers.

You can configure the range for the lockout period for DHCPv4 relay, DHCPv6 relay, DHCPv4 local server, and DHCPv6 local server. You can configure the period globally for all relay agent or local server interfaces, for a group of interfaces, or for specific interfaces within a group. For DHCPv4 relay and local server, you can also configure the lockout for a dual-stack group.

When you enable short-cycle protection, you must specify both the minimum and the maximum duration of the lockout period.

To configure the lockout range for DHCPv4 relay agent:

- Specify the minimum and maximum lockout times.
 - For all DHCPv4 relay agents:

```
[edit forwarding-options dhcp-relay]
user@host# set short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific group of DHCPv4 relay interfaces:

```
[edit forwarding-options dhcp-relay]
user@host# set group group-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific interface within a specified group of DHCPv4 relay interfaces:

```
[edit forwarding-options dhcp-relay]
user@host# set group group-name interface interface-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a DHCPv4 relay dual-stack group:

```
[edit forwarding-options dhcp-relay]
user@host# set dual-stack-group dual-stack-group-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

To configure the lockout range for DHCPv6 relay agent:

- Specify the minimum and maximum lockout times.
- For all DHCPv6 relay agents:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific group of DHCPv6 relay interfaces:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set group group-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific interface within a specified group of DHCPv6 relay interfaces:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set group group-name interface interface-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

To configure the lockout range for DHCPv4 local server:

- Specify the minimum and maximum lockout times.
- For all DHCPv4 local servers:

```
[edit system services dhcp-local-server]
user@host# set short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific group of DHCPv4 local server interfaces:

```
[edit system services dhcp-local-server]
user@host# set group group-name short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific interface within a specified group of DHCPv4 local server interfaces:

```
[edit system services dhcp-local-server]
user@host# set group group-name interface interface-name short-cycle-protection <lockout-
max-time seconds> <lockout-min-time seconds>
```

- For a DHCPv4 local server dual-stack group:

```
[edit system services dhcp-local-server]
user@host# set dual-stack-group dual-stack-group-name short-cycle-protection <lockout-max-
time seconds> <lockout-min-time seconds>
```

To configure the lockout range for DHCPv6 local server:

- Specify the minimum and maximum lockout times.
- For all DHCPv6 local servers:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>
```

- For a specific group of DHCPv6 local server interfaces:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set group group-name short-cycle-protection <lockout-max-time seconds> <lockout-
min-time seconds>
```

- For a specific interface within a specified group of DHCPv6 local server interfaces:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set group group-name interface interface-name short-cycle-protection <lockout-
max-time seconds> <lockout-min-time seconds>
```

Verifying and Managing DHCP Short-Cycle Protection

IN THIS SECTION

- Purpose | [1222](#)
- Action | [1222](#)
- Meaning | [1223](#)

Purpose

View or clear information about DHCP short-cycle protection operations.

Use the supported `show` and `clear` commands to manage and display information about the short-cycle protection operations for the DHCP relay agent and the DHCP local server. You can display information about all locked-out entries or about only individual entries identified by their database index number.

Action

- To display short-cycle protection information for DHCPv4 or DHCPv6 relay agent:

```
user@host> show dhcp relay lockout-entries (all | index index)
user@host> show dhcpv6 relay lockout-entries (all | index index)
```

- To clear short-cycle protection information for DHCPv4 or DHCPv6 relay agent:

```
user@host> clear dhcp relay lockout-entries (all | index index)
user@host> clear dhcpv6 relay lockout-entries (all | index index)
```

- To display short-cycle protection information for DHCPv4 or DHCPv6 local server:

```
user@host> show dhcp server lockout-entries (all | index index)
user@host> show dhcpv6 server lockout-entries (all | index index)
```

- To clear short-cycle protection information for DHCPv4 or DHCPv6 local server:

```
user@host> clear dhcp server lockout-entries (all | index index)
user@host> clear dhcpv6 server lockout-entries (all | index index)
```

Meaning

When you include the `all` option with these `show` commands, information is provided for each client entry in the lockout database, such as the index number that corresponds to the entry in the database, the client identification key, the state of the lockout, how many seconds until the current state is over, how long the current state has been in effect, and how many consecutive times the client has been locked out.

When you want to remove information from the lockout database for a particular client, you must first issue the corresponding `show` command with the `all` option to determine the index for the client entry. Then you can specify that index with the `clear` command.

In the following example, you display all locked-out client entries for DHCPv4 relay agent to find the index number for a particular client, then you clear only that entry and verify that it is deleted:

```
user@host> show dhcp relay lockout-entries all
```

Index	Key	State	Expires(s)	Elapsed(s)	Count
1	00:00:5E:00:53:00	LT	30	5200	2
2	00:00:5E:00:53:11	GT	120	780	2
3	00:00:5E:00:53:22	LT	180	2300	1

```
user@host> clear dhcp relay lockout-entries index 2
```

```
user@host> show dhcp relay lockout-entries all
```

Index	Key	State	Expires(s)	Elapsed(s)	Count
1	00:00:5E:00:53:00	LT	30	5200	2
3	00:00:5E:00:53:22	LT	180	2300	1

In the following example, you display all locked-out client entries for DHCPv6 local server, then you clear all entries and verify that they are deleted:

```
user@host> show dhcp relay lockout-entries all
```

Index	Key	State	Expires(s)	Elapsed(s)	Count
1	00:00:5E:00:53:00	LT	30	5200	2
2	00:00:5E:00:53:11	GT	120	780	2
3	00:00:5E:00:53:22	LT	180	2300	1

```
user@host> clear dhcp relay lockout-entries all
```

```
user@host> show dhcp relay lockout-entries all
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.2R1	Starting in Junos OS Release 18.2R1, <i>DHCP short cycle protection</i> , also called <i>DHCP client lockout</i> , enables the router to reduce these loads by identifying and temporarily locking out clients that continually fail negotiation and have short negotiation cycles as well as clients that frequently complete connections but log out soon after logging in.

RELATED DOCUMENTATION

[DHCP Overview](#) | [1015](#)

[DHCPv6 Local Server](#) | [1449](#)

[DHCPv6 Relay Agent](#) | [1455](#)

DHCP Monitoring and Management

IN THIS SECTION

- [Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings](#) | [1225](#)
- [Viewing and Clearing DHCP Bindings](#) | [1226](#)
- [Monitoring DHCP Relay Server Responsiveness](#) | [1228](#)
- [Verifying DHCP Server Binding and Server Statistics](#) | [1229](#)
- [Verifying and Managing DHCP Relay Configuration](#) | [1230](#)
- [Tracing Extended DHCP Operations](#) | [1231](#)

Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings

You can request that the DHCP local server initiate reconfiguration of all of clients or only specified clients.

To request reconfiguration of all clients:

- Specify the `all` option.

```
user@host> request dhcp server reconfigure all
```

You can use any of the following methods to request reconfiguration of specific clients:

- Specify the IP address of the DHCPv4 client.

```
user@host> request dhcp server reconfigure 192.168.27.3
```

- Specify the MAC address of a DHCPv4 client.

```
user@host> request dhcp server reconfigure 00:00:5E:00:53:67
```

- Specify an interface; reconfiguration is attempted for all clients on this interface.

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

- Specify a logical system; reconfiguration is attempted for all clients or the specified clients in this logical system.

```
user@host> request dhcp server reconfigure all logical-system ls-bldg5
```

- Specify a routing instance; reconfiguration is attempted for all clients or the specified clients in this routing instance.

```
user@host> request dhcp server reconfigure all routing-instance ri-boston
```


Viewing and Clearing DHCP Bindings

This topic provides the procedure you use to display current DHCP bindings, clear selected bindings, and verify that the specified bindings are successfully cleared.

Subscriber management enables you to clear DHCP bindings at several different levels for DHCP local server and DHCP relay agent. For example, you can clear the DHCP bindings on all interfaces, a group of interfaces, or a specific interface. You can also clear DHCP bindings based on IP address, MAC address, session-ID, DHCPv6 prefix, DHCPv6 Client ID, FPC, PIC, port, VLAN, or stacked VLAN (S-VLAN).

This topic includes examples to show several variations of the clear DHCP binding feature. The examples use DHCP local server commands; however, the procedure and commands are similar for DHCP relay agent, DHCPv6 local server, and DHCPv6 relay agent.

To clear bindings and verify the results for a specific IP address:

1. Display current bindings. Issue the appropriate variation of the `show dhcp server binding` command.

```
user@host> show dhcp server binding
2 clients, (2 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
192.168.32.1	00:00:5E:00:53:01	active	2011-10-17 11:38:47 PST
192.168.32.3	00:00:5E:00:53:02	active	2011-00-17 11:38:41 PST

2. Clear the binding you want to remove.

```
user@host> clear dhcp server binding 192.168.32.1
```

3. Verify that the binding has been cleared.

```
user@host> show dhcp server binding
1 clients, (1 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
192.168.32.3	00:00:5E:00:53:01	active	2011-00-17 11:38:41 PST

The following examples show variations of the clear DHCP binding feature. The examples use the DHCP local server version of the commands.



NOTE: IP demux interfaces are not supported by the `show` and `clear` DHCP bindings commands for DHCP local server and DHCP relay agent.

To clear all bindings:

```
user@host> clear dhcp server binding all
```

To clear bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

To clear all bindings over an interface. This example uses the wildcard option.

```
user@host> clear dhcp server binding ge-1/0/0. *
```

To clear bindings on top of a specific VLAN. This example clears all bindings on top of VLAN 100.

```
user@host> clear dhcp server binding ge-1/0/0:100
```

To clear bindings for a specific S-VLAN. This example clears bindings on S-VLAN 100-200.

```
user@host> clear dhcp server binding ge-1/0/0:100-200
```

To clear all bindings on top of all demux VLANs:

```
user@host> clear dhcp server binding demux0
```

To clear all bindings on top of an underlying interface. This example clears the bindings on all demux VLANs on top of interface `ae0`:

```
user@host> clear dhcp server binding ae0
```

To clear PPP bindings. This example uses the wildcard feature and clears the PPP bindings over interface `pp0.100` and `pp0.200`.

```
user@host> clear dhcp server binding pp0.*
```

To clear all bindings on an FPC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1.

```
user@host> clear dhcp server binding ge-1/*
```

To clear all bindings on a PIC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0.

```
user@host> clear dhcp server binding ge-1/0/*
```

To clear all bindings on a port. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0, port 0.

```
user@host> clear dhcp server binding ge-1/0/0.*
```

Monitoring DHCP Relay Server Responsiveness

You can configure DHCP relay agent and DHCPv6 relay agent to enable the router to monitor DHCP server responsiveness. To monitor DHCP server responsiveness, you specify the length of time during which the router tracks how DHCP servers respond to relayed packets. If a configured DHCP server within the routing instance fails to respond to all relayed packets during the specified time period, the router generates the `DH_SVC_EXTERN_SERVER_STATE_CHG` system log message. When the DHCP server begins responding successfully, the router generates the log message again to indicate that responsiveness is restored. You can also use `show dhcp relay statistics` and `show dhcpv6 relay statistics` commands to display DHCP server responsiveness statistics.

The following procedure describes how to configure DHCP relay agent to enable the router to monitor DHCP server responsiveness. To configure DHCPv6 server responsiveness, include the `server-response-time` statement at the `[edit forwarding-options dhcp-relay]` hierarchy level.

To monitor DHCP server responsiveness:

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

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

- 2.

```
[edit forwarding-options dhcp-relay]
user@host# set server-response-time 86,400
```

Verifying DHCP Server Binding and Server Statistics

IN THIS SECTION

- [Purpose | 1229](#)
- [Action | 1229](#)

Purpose

View or clear information about client address bindings and statistics for the extended DHCP local server.



NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Action

- To display the address bindings in the client table on the extended DHCP local server:

```
user@host> show dhcp server binding
```

- To display extended DHCP local server statistics:

```
user@host> show dhcp server statistics routing-instance customer routing instance
```

- To display the address bindings in the client table on the extended DHCP local server at routing-instance level:

```
user@host> show dhcp server binding routing-instance customer routing instance
```

- To display extended DHCP local server statistics at routing-instance level:

```
user@host> show dhcp server statistics routing-instance customer routing instance
```

- To clear the binding state of a DHCP client from the client table on the extended DHCP local server at routing-instance level:

```
user@host> clear dhcp server binding routing-instance customer routing instance
```

- To clear all extended DHCP local server statistics:

```
user@host> clear dhcp server statistics
```

- To clear the binding state of a DHCP client from the client table on the extended DHCP local server:

```
user@host> clear dhcp server binding
```

- To clear all extended DHCP local server statistics at routing-instance level:

```
user@host> clear dhcp server statistics routing-instance customer routing instance
```

Verifying and Managing DHCP Relay Configuration

IN THIS SECTION

- [Purpose | 1231](#)
- [Action | 1231](#)

Purpose

View or clear address bindings or statistics for extended DHCP relay agent clients:

Action

- To display the address bindings for extended DHCP relay agent clients:

```
user@host> show dhcp relay binding routing-instance customer routing instance
```

- To display extended DHCP relay agent statistics:

```
user@host> show dhcp relay statistics routing-instance customer routing instance
```

- To clear the binding state of DHCP relay agent clients:

```
user@host> clear dhcp relay binding routing-instance customer routing instance
```

- To clear all extended DHCP relay agent statistics:

```
user@host> clear dhcp relay statistics routing-instance customer routing instance
```

Tracing Extended DHCP Operations

IN THIS SECTION

- [Configuring the Extended DHCP Log Filename | 1233](#)
- [Configuring the Number and Size of Extended DHCP Log Files | 1233](#)
- [Configuring Access to the Extended DHCP Log File | 1234](#)
- [Configuring a Regular Expression for Extended DHCP Messages to Be Logged | 1235](#)
- [Configuring the Extended DHCP Tracing Flags | 1236](#)
- [Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged | 1236](#)
- [Tracing Extended DHCP Operations for Specific Interfaces | 1237](#)

Both the extended DHCP local server and the extended DHCP relay agent support tracing operations. DHCP tracing operations track extended DHCP operations and record them in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

You can configure DHCP trace operations at the global level and at the interface level. Global DHCP tracing logs all DHCP-related events, whereas interface-level tracing logs only interface-specific DHCP events. If you configure interface-level trace operations, you can specify tracing for a range of interfaces or an individual interface. However, only a single interface-level log file is supported. That is, you cannot specify different interface-level log files for different interfaces or groups of interfaces.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

- Important events for both global and per-interface tracing are logged in a file located in the `/var/log` directory. By default, the router uses the filename, `jdhcpd`. You can specify a different filename, but you cannot change the directory in which trace files are located.
- 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 global DHCP tracing operations.

- Specify tracing operations for DHCP local server and DHCP relay:

```
[edit system processes dhcp-service]
user@host# edit traceoptions
```

The tracing configuration is applied globally to all DHCP applications in every LS:RI. Configuration of event tracing on a per-LS:RI basis is not supported. DHCP tracing is configurable only in the default LS:RI. However, DHCP applications (local server or relay) do not have to be configured in the default LS:RI.



NOTE: We recommend that you use configure tracing statements at the `[edit system processes dhcp-service]` hierarchy level.

Because you can configure DHCP tracing at three different hierarchy levels (one new and recommended, two old and deprecated), the following rules apply to manage the interaction:

- When you configure a filename or any other options for the trace log file, the configuration at the [edit system processes dhcp-service] hierarchy level has the highest precedence, followed by the configuration at the [edit system services dhcp-local-server] hierarchy level, and finally with the lowest precedence, the configuration at the [edit forwarding-options dhcp-relay] hierarchy level.
- The flag configurations for multiple hierarchy levels are merged and applied to all trace log events.
- The deprecated statements do not support filtering the generation of DHCP trace log events by severity level. If you use these statements, trace logging operates with an implicit severity of all, regardless of the severity level configured at the [edit system processes dhcp-service] hierarchy level.

For information about configuring per-interface tracing options, see ["Tracing Extended DHCP Operations for Specific Interfaces" on page 1231](#).

The extended DHCP traceoptions operations are described in the following sections:

Configuring the Extended DHCP Log Filename

By default, the name of the file that records trace output is `jdhcpd`. You can specify a different name by including the `file` option. DHCP local server and DHCP relay agent both support the `file` option for the `traceoptions` statement and the `interface-traceoptions` statement.

To change the filename:

- Specify a filename for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename
```

- Specify a filename for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename
```

Configuring the Number and Size of Extended DHCP 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*.

DHCP local server and DHCP relay agent both support the files and size options for the `traceoptions` statement and the `interface-traceoptions` statement. To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename files number size maximum-file-size
```

- Specify the name, number, and size of the file used for the trace output for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename files number size maximum-file-size
```

Configuring Access to the Extended DHCP 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.

DHCP local server and DHCP relay agent both support the `world-readable` option and the `no-world-readable` option for the `traceoptions` statement and the `interface-traceoptions` statement. To specify that all users can read the log file:

- Configure the log file to be world-readable for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename world-readable
```

- Configure the log file to be world-readable for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename world-readable
```

To explicitly set the default behavior, in which the log file can only be read by the user who configured tracing:

- Configure the log file to be no-world-readable for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename no-world-readable
```

- Configure the log file to be no-world-readable for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename no-world-readable
```

Configuring a Regular Expression for Extended DHCP 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.

DHCP local server and DHCP relay agent both support the `match` option for the `traceoptions` statement and the `interface-traceoptions` statement. To configure regular expressions to be matched:

- Specify the regular expression for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename match regular-expression
```

- Specify the regular expression for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename match regular-expression
```

Configuring the Extended DHCP 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.

DHCP local server and DHCP relay agent both support the `flag` option for the `traceoptions` statement and the `interface-traceoptions` statement. A smaller set of flags is supported for interface-level tracing than for global tracing. To configure the flags for the events to be logged:

- Specify the flags for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set flag flag
```

- Specify the flags for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set flag flag
```

Configuring the Severity Level to Filter Which Extended DHCP 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`.

DHCP local server and DHCP relay agent both support the `level` option for the `traceoptions` statement and the `interface-traceoptions` statement. To configure the flags for the events to be logged:

- Specify the severity level for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set level severity
```

- Specify the severity level for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set level severity
```

Tracing Extended DHCP Operations for Specific Interfaces

In addition to the global DHCP tracing operations, subscriber management enables you to trace extended DHCP operations for a specific interface or for a range of interfaces.

Configuring per-interface tracing is a two-step procedure. In the first step, you specify the tracing options that you want to use, such as file information and flags. In the second step, you enable the tracing operation on the specific interfaces.

To configure per-interface tracing operations:

1. Specify the tracing options you want to use.



NOTE: Per-interface tracing uses the same default tracing behavior as the global extended DHCP tracing operation. The default behavior is described in "[Tracing Extended DHCP Operations](#)" on page 1231.

- a. Specify that you want to configure per-interface tracing options.
 - For DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent:

```
[edit system processes dhcp-service]
user@host# edit interface-traceoptions
```

b. (Optional) Specify the tracing file options.

- Configure the name for the file used for the trace output.

See ["Configuring the Extended DHCP Log Filename" on page 1231.](#)

- Configure the number and size of the log files.

See ["Configuring the Number and Size of Extended DHCP Log Files" on page 1231.](#)

- Configure access to the log file.

See ["Configuring Access to the Extended DHCP Log File" on page 1231.](#)

- Configure a regular expression to filter logging events.

See ["Configuring a Regular Expression for Extended DHCP Messages to Be Logged" on page 1231.](#)

c. (Optional) Specify tracing flag options.

See ["Configuring the Extended DHCP Tracing Flags" on page 1231.](#)

d. (Optional) Configure a severity level for messages to specify which event messages are logged.

See ["Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged" on page 1231.](#)

2. Enable tracing on an interface or interface range.

The following examples show a DHCP local server configuration. You can also use the trace statement at the [edit forwarding-options dhcp-relay] hierarchy level and at the [edit system services dhcp-local-server dhcpv6] hierarchy level.

- Enable tracing on a specific interface.

```
[edit system services dhcp-local-server]
user@host# set group group-name interface interface-name trace
```

- Enable tracing on a range of interfaces.

```
[edit system services dhcp-local-server]
user@host# set group group-name interface interface-name upto interface interface-name
trace
```

RELATED DOCUMENTATION

Dynamic Reconfiguration of Clients From a DHCP Local Server 1197
DHCPv6 Monitoring and Management 1464
Conserving IP Addresses Using DHCP Auto Logout 1208
DHCP Overview 1015
DHCPv6 Local Server 1449
DHCPv6 Relay Agent 1455

3

PART

IPv6 for Subscriber Management

- [IPv6 for Subscriber Management](#) | **1241**
-

IPv6 for Subscriber Management

IN THIS CHAPTER

- Introduction to IPv6 Addresses | 1241
- Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245
- IPv6 WAN Link Addressing with NDRA | 1250
- IPv6 WAN Link Addressing with DHCPv6 IA_NA | 1257
- Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 1259
- WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1269
- Designs for IPv6 Addressing in a Subscriber Access Network | 1305
- Dual-Stack Access Models in a DHCP Network | 1312
- Dual-Stack Access Models in a PPPoE Network | 1324
- Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1351
- Dual Stack for PPPoE Access Networks Using DHCP | 1356
- Dual Stack for PPPoE Access Networks Using NDRA | 1360
- IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services | 1422
- Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | 1425
- Dual Stack Subscribers Monitoring and Management | 1436

Introduction to IPv6 Addresses

IN THIS SECTION

- IPv6 Notation | 1242
- IPv6 Prefixes | 1242
- IPv6 Address Types | 1243

IPv6 uses a 128-bit addressing model compared with the 32-bit addresses used for IPv4. In addition to being larger, IPv6 addresses differ from IPv4 addresses in several ways:

- Notation
- Prefixes
- Address types

These differences give IPv6 addressing greater simplicity and scalability than IPv4 addressing gives.

IPv6 Notation

IPv6 addresses are 128 bits long (expressed as 32 hexadecimal numbers) and consist of eight colon-delimited sections. Each section contains 2 bytes, and each byte is expressed as a hexadecimal number from 0 through FF.

An IPv6 address looks like this:

2001:0db8:0000:0000:0000:0800:200c:7334

By omitting the leading zeroes from each section or substituting contiguous sections that contain zeroes with a double colon, you can write the example address as:

2001:db8:0:0:0:800:200c:7334 or 2001:db8::800:200c:7334

You can use the double-colon delimiter only once within a single IPv6 address. For example, you cannot express the IPv6 address 2001:db8:0000:0000:ea34:0000:71ff:fe01 as 2001:db8::ea34::71ff:fe01.

IPv6 Prefixes

An IPv6 address prefix represents a block of address space or a network. The prefix is a combination of an IPv6 prefix (address) and a prefix length. It takes the form *ipv6-prefix/prefix-length*.

IPv6 addresses can be broken into prefixes of varying length. The prefix length is a decimal value that specifies the number of the leftmost bits in the address that make up the prefix. The prefix length follows a forward slash and, in most cases, identifies the portion of the address owned by an organization. All remaining bits (up to the right-most bit) represent individual nodes or interfaces.

For example, 2001:db8:0000:0000:250:af:34ff:fe26/64 has a prefix length of 64.

The first 64 bits of this address (2001:db8:0000:0000) are the prefix. The rest (250:af:34ff:fe26) identify the interface.

IPv6 Address Types

IN THIS SECTION

- [Unicast Addresses | 1243](#)
- [Multicast Addresses | 1244](#)
- [Anycast Addresses | 1244](#)

There are three major categories of IPv6 addresses:

- Unicast—For a single interface.
- Multicast—For a set of interfaces on the same physical medium. A packet is sent to all interfaces associated with the address.
- Anycast—For a set of interfaces on different physical media. A packet is sent to only one of the interfaces associated with this address, not to all the interfaces.

Unicast Addresses

A unicast address identifies a single interface. When a network device sends a packet to a unicast address, the packet goes only to the specific interface identified by that address. Unicast addresses support a global address scope and two types of local address scopes.

A unicast address consists of n bits for the prefix, and $128 - n$ bits for the interface ID.

In the IPv6 implementation for a subscriber access network, the following types of unicast addresses can be used:

- Global unicast address—A unique IPv6 address assigned to a host interface. These addresses have a global scope and essentially the same purposes as IPv4 public addresses. Global unicast addresses are routable on the Internet.
- Link-local IPv6 address—An IPv6 address that allows communication between neighboring hosts that reside on the same link. Link-local addresses have a local scope, and cannot be used outside the link. They always have the prefix FE80::/10.
- Loopback IPv6 address—An IPv6 address used on a loopback interfaces. The IPv6 loopback address is 0:0:0:0:0:0:0:1, which can be notated as ::1/128.
- Unspecified address—An IPv6 unspecified address is 0:0:0:0:0:0:0:0, which can be notated as ::/128.

Multicast Addresses

A multicast address identifies a set of interfaces that typically belong to different nodes. When a network device sends a packet to a multicast address, the device broadcasts the packet to all interfaces identified by that address. IPv6 does not support broadcast addresses, but instead uses multicast addresses in this role.

Multicast addresses support 16 different types of address scope, including node, link, site, organization, and global scope. A 4-bit field in the prefix identifies the address scope.

The following types of multicast addresses can be used in an IPv6 subscriber access network:

- Solicited-node multicast address—Neighbor Solicitation (NS) messages are sent to this address.
- All-nodes multicast address—Router Advertisement (RA) messages are sent to this address.
- All-routers multicast address—Router Solicitation (RS) messages are sent to this address.

Multicast addresses use the prefix FF00::/8.

Anycast Addresses

An anycast address identifies a set of interfaces that typically belong to different nodes. Anycast addresses are similar to multicast addresses, except that packets are sent only to one interface, not to all interfaces. The routing protocol used in the network usually determines which interface is physically closest within the set of anycast addresses and routes the packet along the shortest path to its destination.

There is no difference between anycast addresses and unicast addresses except for the subnet-router address. For an anycast subnet-router address, the low-order bits, typically 64 or more, are zero. Anycast addresses are taken from the unicast address space.

For more information about anycast addresses, see RFC 2526, *Reserved IPv6 Subnet Anycast Addresses*.

RELATED DOCUMENTATION

| [IPv6 Addressing Requirements for a Subscriber Access Network](#) | 1248

Migration to IPv6 Using IPv4 and IPv6 Dual Stack

IN THIS SECTION

- [Basic Architecture of a Subscriber Access Dual-Stack Network | 1245](#)
- [Terms Used in IPv6 Subscriber Management Documentation | 1246](#)
- [IPv6 Addressing Requirements for a Subscriber Access Network | 1248](#)


As a service provider, you can use the Junos OS IPv4/IPv6 dual-stack feature to begin your migration from IPv4 to IPv6 by implementing IPv6 alongside IPv4 in your existing subscriber networks. The feature allows you to implement IPv6 so that you can provide the same subscriber services over IPv6—video, voice, high-quality data—that you currently provide in your IPv4 networks. You can then perform incremental upgrades to IPv6 and avoid service disruptions while migrating from IPv4 to IPv6.

Basic Architecture of a Subscriber Access Dual-Stack Network

This Juniper Networks dual-stack architecture is designed for either DHCP-based or PPP/PPPoE-based subscriber access networks. In addition, this design allows you to layer DHCPv6 over PPPoE-based networks.

[Figure 8 on page 1246](#) shows the components of a basic subscriber access network in which the subscriber LAN is running both IPv4 and IPv6 and is connected to the IPv4 and IPv6 core using a broadband network gateway (BNG) configured for dual stack. Using IPv4/IPv6 dual stack, the BNG can provide both IPv4 and IPv6 services over the access network to the subscriber LAN. A single interface can operate simultaneously in IPv4 and IPv6 modes.

Figure 8: IPv4 and IPv6 Dual-Stack Architecture in a Subscriber Access Network



Terms Used in IPv6 Subscriber Management Documentation

[Table 54 on page 1246](#) defines terms used in the IPv6 subscriber management documentation.

Table 54: IPv6 Subscriber Management Terms

Term	Definition
BNG	Broadband network gateway. An IP edge router in which bandwidth and QoS policies may be applied. The BNG may encompass any or all of the functionality of B-RAS.
CPE	Customer premises equipment on the subscriber network that connects the subscriber network to the BNG.
Delegated addressing	Method of address assignment in which a host uses IPv6 prefixes to delegate IPv6 global addresses. In a dual-stack network, the CPE uses IPv6 prefixes that it receives to delegate global IPv6 addresses to individual subscriber equipment.
Delegating router	Role of the BNG when it delegates IPv6 prefixes to the requesting router (the CPE).

Table 54: IPv6 Subscriber Management Terms (*Continued*)

Term	Definition
DHCPv6 IA	<p>Identity association. A collection of addresses assigned to a client.</p> <p>Each IA contains one type of address. For example, IA_NA carries assigned addresses that are nontemporary addresses; IA_PD carries a prefix.</p>
DHCPv6 IA_PD	<p>IA for prefix delegation. An IA that carries a prefix that is assigned to the requesting router. Instead of assigning a single address, IA_PD assigns a prefix or a complete subnet.</p> <p>Referred to as DHCPv6 prefix delegation.</p>
DHCPv6 IA_NA	<p>IA for nontemporary addresses. An IA that carries assigned addresses that are not temporary addresses.</p> <p>DHCPv6 IA_NA is used to assign global IPv6 addresses.</p>
Global IPv6 address	<p>Unique IPv6 address that identifies a single interface and allows the interface to access the IPv6 internet.</p>
IPv6 address prefix/ prefix length	<p>Combination of an IPv6 prefix (address) and a prefix length.</p> <p>The prefix takes the form <i>ipv6-prefix/prefix-length</i> and represents a block of address space (or a network).</p> <p>The <i>/prefix-length</i> indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address.</p> <p>For example, 2001:DB8::/32 is an IPv6 prefix.</p>
IPCP	<p>IPv4 Control Protocol. A PPP protocol that establishes the IPv4 link between the BNG and the CPE if you are using PPPoE in your access network.</p>
IPv6CP	<p>IPv6 Control Protocol. A PPP protocol that establishes the IPv6 link between the BNG and the CPE if you are using PPPoE in your access network.</p>

Table 54: IPv6 Subscriber Management Terms (Continued)

Term	Definition
Link-local address	<p>Locally derived address that is designed to be used for addressing on a single link for purposes such as automatic address configuration, Neighbor Discovery, or when no routers are present. It is indicated by the prefix FE80::/10.</p> <p>In your dual-stack network, you can use a link-local address on the interface that connects the CPE and the BNG.</p>
NDRA	Neighbor Discovery Router Advertisement. An IPv6 protocol that is used in the dual-stack network to allow automatic addressing on the CPE WAN link.
Neighbor discovery	Protocol in the IPv6 protocol suite that allows nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors.
Prefix list	Table that contains IPv6 prefixes.
Requesting router	Role of the CPE when it requests IPv6 prefixes from the delegating router (the BNG).
Router Advertisement (RA)	<p>Message that the BNG periodically sends to hosts or sends in response to Router Solicitation (RS) requests from another host. The message includes IPv6 prefixes and other autoconfiguration information.</p> <p>In a dual-stack network, the router sends RAs to CPE devices on its access network.</p>
Router Solicitation (RS)	Message that hosts send to discover the presence of on-link routers. In a dual-stack network, CPE devices send RS messages to the BNG.
Unnumbered address	Address that can be used on the router's PPPoE loopback interface that connects to the CPE.

IPv6 Addressing Requirements for a Subscriber Access Network

IN THIS SECTION


- [Alternatives to Using a Global IPv6 Address on the CPE WAN Link | 1249](#)

You need to implement two types of addressing for IPv6 in a subscriber access network:

- WAN link addressing—For the WAN interface on the CPE (CPE upstream interface).
- Subscriber LAN addressing—For devices connected to the CPE on the subscriber LAN (CPE downstream interfaces).

Figure 9 on page 1249 shows where WAN link addressing and subscriber addressing are assigned in a dual-stack network.

Figure 9: IPv6 Address Requirements in a Subscriber Access Network



g017542

You can use the following methods for assigning IPv6 addresses:

- For WAN link addressing, you can use Neighbor Discovery Router Advertisement (NDRA) or DHCPv6 Identity association for nontemporary addresses (IA_NA) to provision a global IPv6 address.
- For subscriber LAN addressing, you can use DHCPv6 prefix delegation to provision global IPv6 addresses to subscribers on the LAN.

Alternatives to Using a Global IPv6 Address on the CPE WAN Link

If the CPE is supplied by or recommended by the service provider, you do not need to provision a unique global IPv6 address on the CPE. In this case, the broadband network gateway (BNG) can use the loopback interface to manage the CPE. You can use one of the following methods to provision an address on the loopback interface:

- Link-local IPv6 address—Can be used on PPPoE access networks. The link-local address is provisioned by appending the interface identifier negotiated by IPv6CP with the IPv6 link-local prefix (FE80::/10).
- Address derived from DHCPv6 prefix delegation—Can be used on PPPoE access networks or on DHCP access networks. If you use DHCPv6 prefix delegation for subscriber addressing, the CPE can use the prefix it receives from the BNG to assign an IPv6 address on the loopback interface between the CPE and the BNG. This address can be used to manage the CPE, and the CPE uses it as a source address when it communicates with the BNG.

RELATED DOCUMENTATION

[IPv6 WAN Link Addressing with NDRA | 1250](#)

[IPv6 WAN Link Addressing with DHCPv6 IA_NA | 1257](#)

[Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 1259](#)

[WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1269](#)

IPv6 WAN Link Addressing with NDRA

IN THIS SECTION

- [Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250](#)
- [IPv6 Neighbor Discovery Protocol Overview | 1252](#)
- [Dynamic Router Advertisement Configuration Overview | 1253](#)
- [Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | 1253](#)
- [Methods for Obtaining IPv6 Prefixes for NDRA | 1255](#)
- [Duplicate Prefix Protection for NDRA | 1256](#)

Using NDRA to Provide IPv6 WAN Link Addressing Overview

In a dual-stack network, NDRA (Neighbor Discovery Router Advertisement) provides a lightweight address assignment method for autoconfiguration of the global IPv6 address on the CPE WAN link. The CPE device can construct its own IPv6 global address by combining the interface ID that is negotiated by IPv6CP and the prefix obtained through NDRA.

Before NDRA can provide IPv6 address information to the CPE, you need to first obtain a link-local address for the CPE WAN link. NDRA provides address assignment in two phases:

1. Link-local address assignment for local connectivity to the BNG
2. Global address assignment for global connectivity

The process is as follows:

1. During IPv6CP negotiation to establish the PPPoE link between the BNG and the CPE, an interface identifier is negotiated for the CPE.
2. The CPE creates a link-local address by appending the interface identifier with the IPv6 link-local prefix (FE80::/10).



NOTE: When the interface ID is 0, such as for Windows 7 clients, PPP uses the subscriber's session ID in place of the interface ID.

The CPE now has IPv6 connectivity to the BNG, and it can use NDRA to obtain its global IPv6 address.

3. The CPE sends a router solicitation message to the BNG.
4. The BNG responds with a router advertisement message that includes an IPv6 prefix with a length of /64.

This prefix can come directly from a local NDRA address pool configured on the BNG.

If you are using AAA, a RADIUS server can specify the prefix in the *Framed-Ipv6-Prefix* attribute, or it can specify an NDRA pool on the BNG from which the prefix is assigned in the *Framed-Ipv6-Pool* attribute.

5. When the CPE receives the 64-bit prefix, it appends its interface ID to the supplied prefix to form a globally routable 128-bit address.
6. The CPE verifies that the global address is unique by sending a neighbor solicitation message destined to the new address. If there is a reply, the address is a duplicate. The process stops and requires operator intervention.

SEE ALSO

[Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation](#) | 1309

[Design 3: IPv6 Addressing with NDRA](#) | 1310

IPv6 Neighbor Discovery Protocol Overview

IN THIS SECTION

- [Neighbor Discovery Messages | 1252](#)

Neighbor Discovery is a protocol in the IPv6 protocol suite that allows nodes on the same link to advertise their existence to their neighbors and to learn about the existence of their neighbors. Neighbor Discovery is built on top of Internet Control Message Protocol version 6 (ICMPv6). It replaces the following IPv4 protocols: Router Discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

Neighbor Discovery uses router advertisement messages to detect neighbors, advertise IPv6 prefixes, assist in address provisioning, and share link parameters such as MTU, hop limit, advertisement intervals, and lifetime.

Neighbor Discovery Messages

Neighbor Discovery uses the following message types:

- Router advertisement (RA)—Messages sent to announce the presence of the router, advertise prefixes, assist in address configuration, and share other link information such as MTU size and hop limit. The IPv6 nodes on the link can use this information to configure themselves with an IPv6 address and routing information such as the default gateway.
- Router solicitation (RS)—Messages sent by IPv6 nodes when they come online to solicit immediate router advertisements from the router. Starting in Junos OS Release 18.1R1, the well-known IPv6 all-routers multicast address, FF02::2, is supported in nondefault routing instances. Without this support, IPv6 router solicitation packets are dropped in nondefault routing instances.
- Neighbor solicitation (NS)—Messages used for duplicate address detection and to test reachability of neighbors.

A host can verify that its address is unique by sending a neighbor solicitation message destined to the new address. If the host receives a neighbor advertisement in reply, the address is a duplicate.

- Neighbor advertisement (NA)—Messages used for duplicate address detection and to test reachability of neighbors. Neighbor advertisements are sent in response to neighbor solicitation messages.

You can specify the information that is sent in router advertisements.

Dynamic Router Advertisement Configuration Overview

In a network deployment where router interfaces are configured statically, you might need to configure the Router Advertisement Protocol on only a small number of interfaces on which it might run. However, in a subscriber access network, static configuration of the Router Advertisement Protocol becomes impractical because the number of interfaces that potentially need the Router Advertisement Protocol increases substantially. In addition, deploying services in a dynamic environment requires dynamic modifications to interfaces as they are created.

Subscriber access supports the configuration of the Router Advertisement Protocol at the [edit dynamic-profiles *profile-name* protocols] hierarchy level. By specifying Router Advertisement Protocol statements within a dynamic profile, you can dynamically apply a Router Advertisement configuration when a subscriber connects to an interface using a particular access technology (for example, DHCP), enabling the subscriber to access a carrier (multicast) network.

To minimally configure the Router Advertisement Protocol requires that you include the router-advertisement statement at the [edit dynamic-profiles *profile-name* protocols] hierarchy level and the interface statement along with the *\$junos-interface-name* dynamic variable. All other statements are optional.



NOTE: Statements used for Router Advertisement Protocol configuration at the [edit dynamic-profiles *profile-name* protocols] hierarchy level are identical in function to those same statements used for static Router Advertisement Protocol configuration, with the exception of the interface and prefix statements, which use dynamic variables.

SEE ALSO

[Dynamic Profiles Overview](#)

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

[Configuring an Address-Assignment Pool Used for Router Advertisements](#) | 1362

Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors

RFC 4861, Neighbor Discovery for IP version 6 (IPv6), defines the Neighbor Discovery protocol, which is used by IPv6 nodes to determine link-layer addresses for neighbors, track reachability of neighbors, and discover routers that can forward packets on behalf of hosts. Routers send router advertisement messages to advertise their presence on the network and their characteristics. Hosts send router solicitation messages to discover routers by requesting that routers respond with router advertisement messages immediately. The router advertisements are sent both periodically (for the life of the interface) and in response to router solicitations received from hosts.

The router sets the interval between all router advertisements at the value specified by the `max-advertisement-interval` statement for the interface that sends the advertisement messages. The default interval is several minutes in duration, 600 seconds, and can be configured up to 1800 seconds.

A shorter interval for the first few advertisements increases the chances that the router is discovered quickly when it first becomes available. Accordingly, for only the first three unsolicited router advertisements, RFC 4861 requires a router to use an interval no greater than 16 seconds. If the router selects a larger interval, the interval is automatically set to 16 seconds for the first three unsolicited router advertisements.

In some customer scenarios, 16 seconds is too large an interval for the initial router advertisements and can result in an unacceptable delay for establishing subscriber sessions. If you want the router to advertise more aggressively for a quicker discovery, you can explicitly configure the `max-advertisement-interval` statement to less than 16 seconds for the interface that sends router advertisements.

However, this statement sets the interval between all advertisements sent on the interface, not just those for the first three unsolicited advertisements. That means that all router advertisement messages are sent at short intervals when you configure a lower range. Some users may find this undesirable, because they prefer to have the router discovered quickly, but once it is known, they want the advertisements to be sent at a slower pace, acting as keepalives for the duration of the interface without generating unnecessary amounts of traffic.

Starting in Junos OS Release 18.2R1, you can configure global override options to set the range from which the router randomly selects an interval for only the initial three router advertisements for all interfaces. Random interval selection reduces the likelihood that messages from one router are synchronized with those of another router. A new random interval value is selected after each advertisement is sent so that the interval varies between successive messages. The range for the interval between subsequent router advertisement messages per dynamic interface is still configured with the `max-advertisement-interval` statement in a dynamic profile.

To configure the interval in a dynamic profile that applies to router advertisement messages on the dynamic interface:

- Configure the interval.

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
user@host# set max-advertisement-interval seconds
```

To configure an interval range for only the initial three advertisement messages on all interfaces:

1. Configure the low end of the interval range.

```
[edit system services subscriber-management overrides]
user@host# set ra-initial-interval-min seconds
```

2. Configure the high end of the interval range.

```
[edit system services subscriber-management overrides]
user@host# set ra-initial-interval-max seconds
```

Consider the following example, where intervals are configured only for router advertisement messages on a dynamic interface. Because the configured interval value is greater than 16, the interval for the first three unsolicited advertisements is always set to 16 seconds. For all subsequent unsolicited advertisements, the router advertisements are sent at an interval of 60 seconds.

```
[edit dynamic-profiles protocols router-advertisement interface $junos-interface-name]
user@host# set max-advertisement-interval 60
```

Now consider the following example, where intervals are configured globally for the first three unsolicited router advertisement messages on all interfaces. All subsequent unsolicited advertisements are configured per dynamic interface.

```
[edit system services subscriber-management overrides]
user@host# set ra-initial-interval-min 3
user@host# set ra-initial-interval-max 9
[edit dynamic-profiles protocols router-advertisement interface $junos-interface-name]
user@host# set max-advertisement-interval 300
```

In this case, the router generates a random interval between 3 seconds and 9 seconds, inclusive, for the first three router advertisement messages on all interfaces. The router sends all subsequent advertisements at an interval of 300 seconds.

Methods for Obtaining IPv6 Prefixes for NDRA

IN THIS SECTION

- [Using AAA RADIUS Server to Obtain IPv6 Prefixes for NDRA | 1256](#)

You can set up the BNG to select IPv6 prefixes used for NDRA through one of the following methods:

- An external source such as a AAA RADIUS server.
- Dynamic assignment from a local pool of NDRA prefixes that is configured on the BNG

Using AAA RADIUS Server to Obtain IPv6 Prefixes for NDRA

When the BNG needs to obtain a prefix for NDRA, it uses the values in one of the following RADIUS attributes that it receives in Access-Accept messages from the RADIUS server:

- *Framed-IPv6-Prefix*—The attribute contains an IPv6 prefix that the BNG can send to the CPE in router advertisement messages.
- *Framed-IPv6-Pool*—The attribute contains the name of an NDRA pool configured on the BNG from which the BNG can select a prefix to include in router advertisements.

SEE ALSO

[Configuring an Address-Assignment Pool Used for Router Advertisements | 1362](#)

Duplicate Prefix Protection for NDRA

If you are using AAA to supply IPv6 prefixes for NDRA, you can enable duplicate prefix protection for NDRA. If enabled, the BNG checks the following attributes received from external servers:

- *Framed-IPv6-Prefix*
- *Framed-IPv6-Pool*

The router then takes one of the following actions:

- If a prefix overlaps with a prefix in an address pool, the prefix is taken from the pool if it is available.
- If the prefix is already in use, it is rejected as unavailable.
- If the prefix length requested from the external server does not match the pool's prefix length exactly, the authentication request is denied. If configured, the Acct-Stop message includes a termination cause.

SEE ALSO

[Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | 1363](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.2R1	Starting in Junos OS Release 18.2R1, you can configure global override options to set the range from which the router randomly selects an interval for only the initial three router advertisements for all interfaces.
18.1R1	Starting in Junos OS Release 18.1R1, the well-known IPv6 all-routers multicast address, FF02::2, is supported in nondefault routing instances. Without this support, IPv6 router solicitation packets are dropped in nondefault routing instances.

RELATED DOCUMENTATION

[Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245](#)

[Designs for IPv6 Addressing in a Subscriber Access Network | 1305](#)

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Dual Stack for PPPoE Access Networks Using NDRA | 1360](#)

IPv6 WAN Link Addressing with DHCPv6 IA_NA

IN THIS SECTION

- [Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA | 1258](#)
- [Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | 1258](#)

You can use DHCPv6 IA_NA to assign a global IPv6 address to the CPE WAN link. If the CPE sends a Solicit message that contains the IA_NA option to the BNG, the BNG acts as a DHCPv6 server and assigns a single IPv6/128 address to the WAN interface of the CPE.

Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA

IN THIS SECTION

- [Using a AAA RADIUS Server to Obtain IPv6 Addresses for DHCPv6 IA_NA | 1258](#)

You can set up the BNG to select global IPv6 addresses to be delegated to the requesting router in one of the following ways:

- An external source such as a AAA RADIUS server or a DHCP server using the DHCPv6 relay agent.
- Dynamic assignment from a local pool of addresses that is configured on the BNG

Using a AAA RADIUS Server to Obtain IPv6 Addresses for DHCPv6 IA_NA

When the BNG needs to obtain a global IPv6 for the CPE WAN link and optionally a DHCPv6 prefix, it uses the values in one of the following RADIUS attributes:

- *Framed-IPv6-Prefix*—The attribute contains a global IPv6 address with a prefix length of 128.
- *Framed-IPv6-Pool*—The attribute contains the name of an address-assignment pool configured on the BNG from which the BNG can select a global IPv6 address to send to the CPE.

Both attributes are sent from the RADIUS server to the BNG in RADIUS Access-Accept messages.

Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA

This procedure shows how to configure IPv6 local address pools to allocate global IPv6 addresses to the CPE WAN link.

You must configure separate pools for DHCPv6 prefix delegation, DHCPv6 IA_NA, and NDRA.

To configure the pool to be used for DHCPv6 IA_NA:

1. Create a pool and assign a name to it.

```
[edit access]
user@host# edit address-assignment pool v6-ia-na-pool
```

2. Under family inet6, add IPv6 network prefixes to the pool.

```
[edit access address-assignment pool v6-ia-na-pool]
user@host# edit family inet6
user@host# set prefix 2001:db8:0000::/64
```

3. Configure the name of the IPv6 address range, and define the range by setting a low and high range of /128 addresses.

```
[edit access address-assignment pool v6-ia-na-pool family inet6]
user@host# edit range v6-range
user@host# set low 2001:db8::1/128
user@host# set high 2001:db8::ffff:ffff/128
```

RELATED DOCUMENTATION

[Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245](#)

[Designs for IPv6 Addressing in a Subscriber Access Network | 1305](#)

[Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1351](#)

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Dual Stack for PPPoE Access Networks Using DHCP | 1356](#)

Subscriber LAN Addressing with DHCPv6 Prefix Delegation

IN THIS SECTION

- [Using DHCPv6 Prefix Delegation Overview | 1260](#)
- [Using a Delegated Prefix on the CPE Loopback Interface | 1261](#)
- [DHCPv6 Prefix Delegation over PPPoE | 1261](#)
- [Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation | 1262](#)
- [Subscriber Delegated Prefix Preservation | 1263](#)
- [DHCPv6 Prefix Exclusion | 1264](#)
- [Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation | 1267](#)

- Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | 1267

Using DHCPv6 Prefix Delegation Overview


You can use DHCPv6 prefix delegation to automate the delegation of IPv6 prefixes to the CPE. With prefix delegation, a delegating router (the BNG) delegates IPv6 prefixes to a requesting router (the CPE). The requesting router then uses the prefixes to assign global IP addresses to the devices on the subscriber LAN. The requesting router can also assign subnet addresses to subnets on the LAN.

DHCPv6 prefix delegation is useful when the delegating router does not have information about the topology of the networks in which the requesting router is located. In such cases, the delegating router requires only the identity of the requesting router to choose a prefix for delegation.

DHCPv6 prefix delegation replaces the need for NAT in an IPv6 network.

[Figure 10 on page 1260](#) shows how DHCPv6 prefix delegation is used in a dual-stack network.

Figure 10: Delegated Addressing in a Dual-Stack Network Using DHCPv6



DHCPv6 prefix delegation operates as follows:

1. A delegating router is provided with IPv6 prefixes to be delegated to requesting routers. These prefixes can come from a local address-assignment pool or an external AAA server.
Each prefix has an associated valid and preferred lifetime, which can be extended.
2. A requesting router requests one or more prefixes from the delegating router.

3. The delegating router chooses prefixes for delegation, and responds with prefixes to the requesting router.
4. The requesting router is then responsible for the delegated prefixes.

The address allocation mechanism in the subscriber network can be performed with ICMPv6 Neighbor Discovery in router advertisements, DHCPv6, or a combination of these two methods.

SEE ALSO

[Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1308](#)

[Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | 1309](#)

[Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | 1311](#)

[Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | 1306](#)

Using a Delegated Prefix on the CPE Loopback Interface

For networks in which the service provider directly controls the CPE, a delegated prefix can be used to create an IPv6 address on the loopback interface between the CPE and the BNG. This address can be used to manage the CPE, and the CPE uses it as a source address when it communicates with the BNG.

SEE ALSO

[Selecting the Type of Addressing Used on the CPE | 1305](#)

DHCPv6 Prefix Delegation over PPPoE

The process of DHCPv6 prefix delegation when DHCPv6 is running over a PPPoE access network is as follows:

1. The CPE obtains a link-local address by appending the interface ID that it receives through IPv6CP negotiation to the IPv6 link-local prefix (FE80::/10). The link-local address provides an initial path for protocol communication between the BNG and CPE
2. The CPE sends a DHCPv6 Solicit message that includes an IA_PD option.
3. The BNG chooses a prefix for the CPE with information from an external AAA server or from a local prefix pool.
4. The BNG sends an Advertise message to the CPE. The message includes the delegated prefix, an IA_PD option, and an IA_PD prefix option. The prefix length in the IA_PD prefix option is 48. The message can also contain other configuration information, such as a maximum lease time.

5. The CPE sends a Request message to the BNG. The message requests the prefix that was advertised.
6. The BNG returns the delegated prefix to the CPE in a Reply message. This message also contains the delegated prefix, an IA_PD option, and an IA_PD prefix option. The prefix length in the IA_PD prefix option is 48. The message can also contain other configuration information, such as a maximum lease time.
7. The CPE uses the delegated prefix to allocate global IPv6 addresses to host devices on the subscriber network. It can use router advertisements, DHCPv6, or a combination of these two methods to allocate addresses on the subscriber LAN.

SEE ALSO

[Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | 1390](#)

Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation

IN THIS SECTION

- [Using a AAA RADIUS Server to Obtain IPv6 Prefixes for Prefix Delegation | 1262](#)

You can set up the BNG to select IPv6 prefixes to be delegated to the requesting router in one the following ways:

- An external source such as a AAA RADIUS server or a DHCP server using the DHCPv6 relay agent.
- Dynamic assignment from a local pool of prefixes that is configured on the BNG

You can specify the name of a delegated pool to use for prefix delegation, which means that you do not need to use AAA to obtain the pool name. In this configuration, if you have also specified a pool match order, the specified delegated pool takes precedence.

Using a AAA RADIUS Server to Obtain IPv6 Prefixes for Prefix Delegation

When the BNG needs to obtain a prefix for DHCPv6 prefix delegation, it uses the values in one of the following RADIUS attributes:

- *Delegated-IPv6-Prefix*—The attribute (123) contains an IPv6 prefix that the BNG can send to the CPE.

- *Jnpr-IPv6-Delegated-Pool-Name*—The attribute (VSA 26-161) contains the name of an address-assignment pool configured on the BNG from which the BNG can select a prefix to send to the CPE.

Both attributes are sent from the RADIUS server to the BNG in RADIUS Access-Accept messages.

SEE ALSO

| [Selecting the Method of Obtaining IPv6 Prefixes](#) | 1307

Subscriber Delegated Prefix Preservation

SUMMARY

Subscriber delegated prefix preservation supports preserving and restoring IPv6 prefixes assigned using DHCPv6 prefix delegation.

Configure address preservation under [edit access] hierarchy:

```
user@host# set access address-preservation address-types delegated-prefix
user@host# set access address-preservation aging-time address-aging-time
```

The subscriber's delegated prefix is preserved even when subscriber logs out. After logout the output of "show network-access address-assignment preserved" command is updated and the subscriber gets allocated the same delegated prefix after login.

```
user@host> show network-access address-assignment preserved
```

Subscriber delegated prefix preservation supports preserving and restoring a subscriber's delegated prefix through multiple logins. Receiving the same delegated prefix from the JUNOS device prevents an IA-PD change, that triggers renegotiation for all hosts on the LAN behind the CPE. To utilise this functionality, a configuration command is executed. This functionality is disabled by default. To enable address preservation for DHCPv6 prefix delegation, use the `access address-preservation` configuration command.

Each participating subscriber is identified by the Agent-Circuit-Id. The Agent-Circuit-Id is available in standard upstream control traffic. (For example: Option-82 sub-option 1 for DHCP, Option 18 for DHCPv6). If the access address-preservation prefix is configured and a subscriber has an Agent-Circuit-Id, the subscribers prefix is preserved through multiple logins. The Agent-Circuit-Id is the key that is

used to store the subscriber's prefix when the subscriber logs out. When the subscriber with the same Agent-Circuit-Id, logs in again the prefix is retrieved and associated to the subscriber.

The `show network-access address-assignment preserved` command supports an optional routing-instance *<routing-instance-name>* output.

```

user@host> show network-access address-assignment preserved routing-instance SUBSCRIBERS
Key                                     Delegated Prefix      Reservation End
interfaceId_Port //1/1                2001:0db8:aaaa:1::/64 May 22 12:06:19

```

SEE ALSO

- [address-preservation \(Access\)](#)
- [show network-access address-assignment pool](#)

DHCPv6 Prefix Exclusion

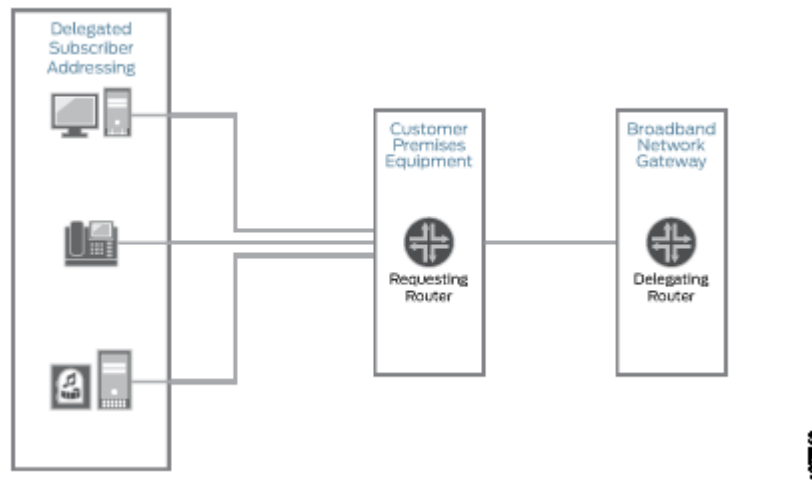
IN THIS SECTION

- [Configuring DHCPv6 Prefix Exclude Option | 1266](#)

You can use the Dynamic Host Configuration Protocol v6 (DHCPv6) prefix delegation to automate the delegation of IPv6 prefixes to the customer premises equipment (CPE) devices. With prefix delegation, a delegating router - the broadband network gateway (BNG) router, delegates IPv6 prefixes to a requesting router such as a CPE device. The requesting router then uses the prefixes to assign global IP addresses to the devices on the subscriber LAN. The requesting router can also assign subnet addresses to subnets on the LAN. DHCPv6 prefix delegation is useful when the delegating router does not have information about the topology of the networks in which the requesting router is located. In such cases, the delegating router requires only the identity of the requesting router to choose a prefix for delegation. DHCPv6 prefix delegation replaces the need for NAT in an IPv6 network.

[Figure 11 on page 1265](#) shows how DHCPv6 prefix delegation is used in a dual-stack network.

Figure 11: Delegated Addressing in a Dual-Stack Network Using DHCPv6



DHCPv6 prefix delegation operates as follows:

1. A delegating router is provided with IPv6 prefixes to be delegated to requesting routers. These prefixes can come from a local address-assignment pool or an external AAA server.
Each prefix has an associated valid and preferred lifetime, which can be extended.
2. A requesting router requests one or more prefixes from the delegating router.
3. The delegating router chooses prefixes for delegation, and responds with prefixes to the requesting router.
4. The requesting router is then responsible for the delegated prefixes.

The address allocation mechanism in the subscriber network can be performed with ICMPv6 Neighbor Discovery Protocol (NDP) in router advertisements, DHCPv6, or a combination of these two methods.

The requesting router cannot use a sub-prefix of the delegated prefix assigned to it by the delegating router to the link between the delegating router and the requesting router. Because of this limitation, there are usually two routes to the CPE device. One is the delegated prefix, for the customer site behind the CPE device and the other for the link between the requesting router and the delegating router. To overcome this, Junos OS allows the exclusion of one specific prefix from a delegated prefix set while using DHCPv6 based prefix delegation as described in RFC 6603. This excluded prefix is used as the link between the delegating router and the requesting router. This prefix link is intended for use in networks where each requesting router is in its own Layer 2 domain.

To support prefix exclude delegation, the requesting router includes the Option Request option (ORO) with the PD_Exclude option in the solicit, request, renew, or rebind message to inform the delegating router about the support for the prefix delegation. When the Juniper Networks router acting as the

DHCP server receives these message and finds the exclude prefix option (option 67) in ORO, it decides the prefix to be excluded. (The length of the prefix to be excluded is bigger than the delegated prefix length.) The excluded prefix is then added in the IA_Prefix options. The DHCP server acting as relay forwards the requested option to the server and relays the excluded prefix, assigned by the server, back to the client.

To exclude a prefix length in a DHCP server, configure the `exclude-prefix-len` statement at the `[edit access address-assignment pool pool-name family dhcpv6 dhcp-attributes]` hierarchy level. The length of the prefix can range from 1 through 128.

If the DHCP server supporting the exclude prefix wants the client to request for a prefix exclude after reconfiguration, then you can configure the `support-option-pd-exclude` statement either at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level or at the `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

Configuring DHCPv6 Prefix Exclude Option

To configure DHCPv6 prefix exclude:

1. Configure the prefix length to be excluded from a delegated prefix set pool. This prefix is used as the link between the delegating router and the requesting router. The exclude prefix length is bigger than the given prefix length.

```
[edit access address-assignment pool pool-name family inet6 dhcp-attributes]
user@host# set exclude-prefix-len prefix-length
```

For example, for prefix delegated in 2001:db8::/32 , configure the exclude prefix as 2001:db8:ffff:fffc::/72 for delegated pool *prefix_delegate_pool*.

```
[edit access address-assignment pool prefix_delegate_pool family inet6 dhcp-attributes]
user@host# set exclude-prefix-len 72
```

2. Configure PD_Exclude option support in the reconfigure message. In case the server wants the client to request for the prefix to be excluded after reconfiguration then the exclude prefix options are added in the Option Request option (ORO) in the reconfigure message.

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set support-option-pd-exclude
```

3. Configure PD_Exclude option support in the reconfigure message for a given group.

```
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
user@host# set support-option-pd-exclude
```

Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation

This procedure shows how to configure IPv6 local address pools to allocate IPv6 prefixes for use by DHCPv6 prefix delegation.

You must configure separate pools for DHCPv6 prefix delegation, DHCPv6 IA_NA, and NDRA.

To configure the pool to be used for prefix delegation:

1. Create a pool and assign a name to it.

```
[edit access]
user@host# edit address-assignment pool v6-prefix-pool-2001
```

2. Under family inet6, add IPv6 prefixes to the pool.

```
[edit access address-assignment pool v6-prefix-pool-2001]
user@host# edit family inet6
user@host# set prefix 2001:db8:0000:0000:0000::/64
```

3. Configure the name of the IPv6 prefix range, and define the range by setting a prefix length of 64.

```
[edit access address-assignment pool v6-prefix-pool-2001 family inet6]
user@host# edit range prefix-range
user@host# set prefix-length 64
```

Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation

You can explicitly specify which address pool the BNG uses to assign IPv6 prefixes for use by DHCPv6 prefix delegation. This feature enables you to identify the address pool without using RADIUS or a network match.



NOTE: If the Juniper Networks IPv6-Delegated-Pool-Name VSA (26–161) provides assigns a delegated address pool, the VSA-specified value takes precedence over the delegated-address statement.



NOTE: You can specify the local delegated address pool at the following levels:

- Globally for the server at the [edit system services dhcp-local-server dhcpv6 overrides] hierarchy level.
- For a named group of interfaces at the [edit system services dhcp-local-server dhcpv6 group *group-name* overrides] hierarchy level.
- For a specific interface within a named group of interface at the [edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides] hierarchy level.

The following steps show only how to specify a local pool used globally by the local server.

To specify the pool to be used for prefix delegation:

1. Specify that you want to configure override options for DHCPv6 local server.

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

2. Specify the name of the delegated address pool.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set delegated-pool pool-name
```

RELATED DOCUMENTATION

[Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245](#)

[Designs for IPv6 Addressing in a Subscriber Access Network | 1305](#)

WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation

IN THIS SECTION

- [Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269](#)
- [DHCPv6 Options in a DHCPv6 Multiple Address Environment | 1271](#)
- [Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA | 1272](#)
- [Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface | 1273](#)
- [Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | 1274](#)

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview

IN THIS SECTION

- [Lease Times and Session Timeouts for DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1269](#)
- [Behavior When CPE Sends Separate Renew Requests for IA_NA and IA_PD Address Types | 1270](#)

You can use DHCPv6 IA_NA to assign a global IPv6 address to the CPE WAN link and DHCPv6 prefix delegation to provide prefixes for use on the subscriber LAN. DHCPv6 IA_NA and DHCPv6 prefix delegation are done in a single DHCPv6 session. If the CPE sends both the IA_NA and IA_PD options in the same DHCPv6 Solicit message, the BNG returns both a single IPv6/128 address and an IPv6 prefix.

When at least one address is successfully allocated, the router creates a subscriber entry and binds the entry to the assigned address. If both addresses are successfully allocated, the router creates a single subscriber entry and binds both addresses to that entry.

Lease Times and Session Timeouts for DHCPv6 IA_NA and DHCPv6 Prefix Delegation

When you use DHCPv6 IA_NA together with DHCPv6 prefix delegation, note the following about session timeouts and lease times:

- A session timeout from AAA has the highest precedence and overrides local pool lease times.

- For DHCPv6 local server, the minimum lease time associated with an address pool takes precedence over pools with longer lease times. For example, if a CPE obtains an IA_NA address from a pool with a lease time of 3600, and a prefix from a pool with a lease time of 7200, the lease time returned in the Reply message from the BNG is 3600.
- If AAA does not return a session timeout and the address pool does not have a configured lease time, the default setting of 86,400 (one day) is used.

Behavior When CPE Sends Separate Renew Requests for IA_NA and IA_PD Address Types

In some networks, the DHCPv6 client CPE device does both of the following:

- Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
- Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

Starting in Junos OS Release 17.2R3, 17.4R2, 18.1R3, 18.2R2, and 18.3R1, the `jdhcpd` process extends the lease for both address types in this situation.

1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases, the behavior is different for this situation:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.



NOTE: For dual-stacked clients over the same session (PPP over L2TP LNS, DHCP, or IPoE), enhanced subscriber management does not support configurations where both of the following are true:

- The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
- The solicit messages specify a type 2 or type 3 DUID (link-layer address).

As a workaround, you must configure the CPE to send a single solicit message for both IA_NA and IA_PD when the other configuration elements are present.

SEE ALSO

[IPv6 WAN Link Addressing with DHCPv6 IA_NA | 1257](#)

[Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1308](#)

DHCPv6 Options in a DHCPv6 Multiple Address Environment

In a DHCPv6 environment, DHCPv6 clients can use a single DHCPv6 Solicit message to request multiple addresses (for example, IA_NA address, IA_PD address, or both), as well as the DNS server address (DHCPv6 attribute 23). When a client requests multiple addresses, DHCPv6 uses the following guidelines to determine how options are returned to the client.

- **DNS server address**—Whenever a client requests an IA_PD address (either alone or with an IA_NA address) and also requests a DNS server address, DHCPv6 returns a DNS address only when one is specified in the IA_PD pool. If the IA_PD pool does not include a DNS address, DHCPv6 ignores any DNS address configured in the IA_NA pool.

If the client requests an IA_NA address (but not an IA_PD address) and also a DNS server address, DHCPv6 returns a DNS address if one is configured in the IA_NA pool.

- **Lease time**—DHCPv6 returns the shortest value of the lease times configured in the IA_NA pool, the IA_PD pool, and authd. DHCPv6 uses this value to set the lifetimes and the Renew and Rebind timers.



NOTE: By default, DHCPv6 local server returns the DNS server address as a global DHCPv6 option. You can override the current default behavior if you want DHCPv6 to return the DNS server address at the suboption level.

SEE ALSO

[Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269](#)

[Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | 1507](#)

Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA

IN THIS SECTION

- [Address Pools for DHCPv6 Prefix Delegation and DHCPv6 IA_NA | 1272](#)
- [Using a AAA RADIUS Server to Obtain IPv6 Addresses and Prefixes | 1272](#)
- [Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment | 1273](#)

You can set up the BNG to select global IPv6 addresses to be delegated to the requesting router in one of the following ways:

- An external source such as a AAA RADIUS server or a DHCP server using the DHCPv6 relay agent.
- Dynamic assignment from a local pool of prefixes or global IPv6 addresses that is configured on the BNG

Address assignment for prefix delegation and IA_NA are independent. For example, you can use AAA RADIUS for DHCPv6 IA_NA, and use a local pool for prefix delegation.

Address Pools for DHCPv6 Prefix Delegation and DHCPv6 IA_NA

You need two separate address pools for prefix delegation and IA_NA. The pool used for IA_NA contains /128 addresses, and the pool for prefix delegation contains /56 or /48 addresses.

You can specify the name of a delegated pool to use for prefix delegation, which means that you do not need to use AAA to obtain the pool name. In this configuration, if you have also specified a pool match order, the specified delegated pool takes precedence.

You can configure pool attributes so that the IA_NA pool and the prefix delegation pool can specify different SIP servers for DNS addresses. DHCPv6 options that the BNG returns to the CPE are based on the pool from which the addresses were allocated. These options that are returned are based on the DHCPv6 Option Request option (ORO), which can be configured globally or within the IA_NA and IA_PD request.

Using a AAA RADIUS Server to Obtain IPv6 Addresses and Prefixes

When the BNG needs to obtain a global IPv6 address for the CPE WAN link and a DHCPv6 prefix, it uses the values in one of the following RADIUS attributes:

- *Framed-IPv6-Prefix*—The attribute contains a global IPv6 address and a prefix. A prefix length of 128 is associated with the global IPv6 address. Prefix lengths less than 128 are associated with prefixes.
- *Framed-IPv6-Pool*—The attribute contains the name of an address-assignment pool configured on the BNG, from which the BNG can select a global IPv6 address or an IPv6 prefix to send to the CPE.

Both attributes are sent from the RADIUS server to the BNG in RADIUS Access-Accept messages.

Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment

To configure dynamic DHCPv6 address assignment for both DHCPv6 IA_NA and DHCPv6 prefix delegation, use the `$junos-subscriber-ipv6-multi-address` predefined variable in your dynamic profile. You use this variable in place of the `$junos-subscriber-ipv6-address` variable, which supports a single IPv6 address or prefix. The `$junos-subscriber-ipv6-multi-address` variable is applied as a demultiplexing source address, and is expanded to include both the host and prefix addresses.

You include the `$junos-subscriber-ipv6-multi-address` variable at the `[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number family inet6 demux-source]` hierarchy level.

SEE ALSO

[Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation | 1267](#)

[Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | 1258](#)

Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface

DHCPv6 relay agent supports multiple IA_NA and IA_PD requests within a single DHCPv6 Solicit message. The requests can be any combination of IA_NA and IA_PD addresses, up to a maximum of eight requests. As part of the multiple IA request support, each address lease is assigned its own lease time expiration, independent of the other leases. The use of independent lease timers ensures that when one lease is torn down, the other active leases are maintained. You can use the `show dhcpv6 relay binding` and `show dhcpv6 relay binding detail` commands to display the status of the individual lease times.

The DHCPv6 support for multiple IA requests enables you to use prefix delegation to designate blocks of addresses, as described in RFC 3633, *IPv6 Prefix Options for DHCPv6*. For example, you might want to delegate multiple address blocks to a customer premises equipment (CPE) router as a means to simplify flow classification and service monetization in your IPv6 environment.

Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE

IN THIS SECTION

- [Requirements | 1274](#)
- [Overview | 1274](#)
- [Configuration | 1276](#)
- [Verification | 1298](#)

Requirements

This example uses the following hardware and software components:

- MX Series 5G Universal Routing Platform
- Junos OS Release 11.4 or later

Overview

IN THIS SECTION

- [Topology | 1275](#)

This design uses DHCPv6 IA_NA and DHCPv6 prefix delegation in your subscriber access network as follows:

- The access network is PPPoE.
- DHCPv6 IA_NA is used to assign a global IPv6 address on the WAN link. The address comes from a local pool that is specified using AAA RADIUS.
- DHCPv6 prefix delegation is used for subscriber LAN addressing. It used a delegated prefix from a local pool that is specified using AAA RADIUS.
- DHCPv4 is used for subscriber LAN addressing.
- DHCPv6 subscriber sessions are layered over an underlying PPPoE subscriber session.

Topology

Figure 12: PPPoE Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

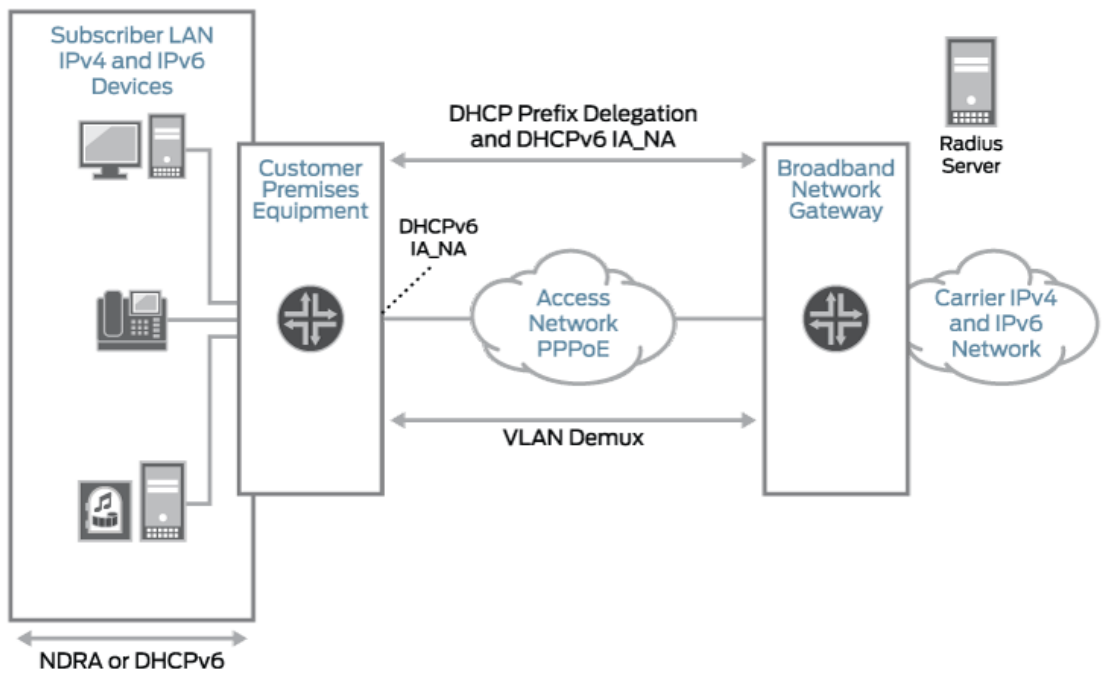


Table 55 on page 1275 describes the configuration components used in this example.

Table 55: Configuration Components Used in Dual Stack with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

Configuration Component	Component Name	Purpose
Dynamic Profile	pppoe-subscriber-profile	Profile that creates a PPPoE logical interface when the subscriber logs in.
Interfaces	ge-0/2/5	Interface used for communication with the RADIUS server.
	ge-0/3/0	Underlying Ethernet interface.

Table 55: Configuration Components Used in Dual Stack with DHCPv6 IA_NA and DHCPv6 Prefix Delegation (Continued)

Configuration Component	Component Name	Purpose
	demux0	VLAN demux interface that runs over the underlying Ethernet interface.
	lo0	Loopback interface for use in the access network. The loopback interface is automatically used for unnumbered interfaces.
Address-Assignment Pools	pool v4-pool	Pool that provides IPv4 addresses for the subscriber LAN.
	pool v6-ia-na-pool	Pool that provides a global IPv6 address to the CPE WAN link.
	pool v6-pd-pool	Pool that provides a pool of prefixes that are delegated to the CPE and used for assigning IPv6 global addresses on the subscriber LAN.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1277](#)
- [Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | 1280](#)
- [Configuring a Dynamic Profile for the PPPoE Logical Interface | 1282](#)
- [Configuring a Loopback Interface | 1285](#)
- [Configuring a VLAN Demux Interface over an Ethernet Underlying Interface | 1287](#)
- [Configuring an Interface for Communication with RADIUS Server | 1289](#)
- [Specifying the BNG IP Address | 1290](#)
- [Configuring RADIUS Server Access | 1292](#)
- [Configuring RADIUS Server Access Profile | 1294](#)

● [Configuring Local Address-Assignment Pools | 1295](#)

CLI Quick Configuration

The following is the complete configuration for this example:

```
dynamic-profiles {
  pppoe-subscriber-profile {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
      }
    }
    interfaces {
      pp0 {
        unit "$junos-interface-unit" {
          ppp-options {
            chap;
            pap;
          }
          pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
          }
          keepalives interval 30;
          family inet {
            unnumbered-address "$junos-loopback-interface";
          }
          family inet6 {
            unnumbered-address "$junos-loopback-interface";
          }
        }
      }
    }
  }
}

system {
  services {
    dhcp-local-server {
```

```

        dhcpv6 {
            group v6-ppp-subscriber {
                interface pp0.0;
            }
        }
    }
}

interfaces {
    ge-0/2/5 {
        gigether-options {
            no-auto-negotiation;
        }
        unit 0 {
            family inet {
                address 203.0.113.99/32;
            }
        }
    }
    ge-0/3/0 {
        hierarchical-scheduler maximum-hierarchy-levels 2;
        flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1;
    }
    demux0 {
        unit 1 {
            proxy-arp;
            vlan-tags outer 1 inner 1;
            demux-options {
                underlying-interface ge-0/3/0;
            }
            family pppoe {
                duplicate-protection;
                dynamic-profile pppoe-subscriber-profile;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 203.0.113.1/32 {
                    primary;

```



```

        }
        dhcp-attributes {
            maximum-lease-time 99999;
        }
    }
}
pool v6-ia-na-pool {
    family inet6 {
        prefix 2001:db8:1000:0000::/64;
        range v6-range-0 {
            low 2001:db8:1000::1/128;
            high 2001:db8:1000::ffff:ffff/128;
        }
    }
}
pool v6-pd-pool {
    family inet6 {
        prefix 2001:db8:2012::/48;
        range v6-pd prefix-length 64;
    }
}
}
address-protection;
}

```

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE

CLI Quick Configuration

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

```

edit system services dhcp-local-server dhcpv6
edit group v6-ppp-subscriber
set interface pp0.0

```

Step-by-Step Procedure

To layer DHCPv6 above the PPPoE IPv6 family (inet6), associate DHCPv6 with the PPPoE interfaces by adding the PPPoE interfaces to the DHCPv6 local server configuration. Because this example uses a dynamic PPPoE interface, we are using the pp0.0 (PPPoE) logical interface as a wildcard to indicate that a DHCPv6 binding can be made on top of a PPPoE interface.

To configure a DHCPv6 local server:

1. Access the DHCPv6 local server configuration.

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

2. Create a group for dynamic PPPoE interfaces and assign a name.

The group feature groups a set of interfaces and then applies a common DHCP configuration to the named interface group.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit group v6-ppp-subscriber
```

3. Add an interface for dynamic PPPoE logical interfaces.

```
[edit system services dhcp-local-server dhcpv6 group v6-ppp-subscriber]
user@host# set interface pp0.0
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit]
user@host# show
system {
  services {
    dhcp-local-server {
      dhcpv6 {
        group v6-ppp-subscriber {
          interface pp0.0;
        }
      }
    }
  }
}
```



```

    }
  }
}

```

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

Configuring a Dynamic Profile for the PPPoE Logical Interface

CLI Quick Configuration

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

```

edit dynamic-profiles pppoe-subscriber-profile
edit routing-instances $junos-routing-instance
set interface $junos-interface-name
exit
edit interfaces pp0 unit $junos-interface-unit
set family inet unnumbered-address "$junos-loopback-interface"
set family inet6 unnumbered-address "$junos-loopback-interface"
set pppoe-options underlying-interface "$junos-underlying-interface"
set pppoe-options server
set ppp-options pap
set ppp-options chap
set keepalives interval 30

```

Step-by-Step Procedure

Create a dynamic profile for the PPPoE logical interface. This dynamic profile supports both IPv4 and IPv6 sessions on the same logical interface.

To configure the dynamic profile:

1. Create and name the dynamic profile.

```

[edit]
user@host# edit dynamic-profiles pppoe-subscriber-profile

```

2. Add a routing instance to the profile.

```
[edit dynamic-profiles pppoe-subscriber-profile]
user@host# edit routing-instances $junos-routing-instance
user@host# set interface $junos-interface-name
```

3. Configure a PPPoE logical interface (pp0) that is used to create logical PPPoE interfaces for the IPv4 and IPv6 subscribers.

```
[edit dynamic-profiles pppoe-subscriber-profile]
user@host# edit interfaces pp0
```

4. Specify \$junos-interface-unit as the predefined variable to represent the logical unit number for the pp0 interface. The variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0]
user@host# edit unit $junos-interface-unit
```

5. Specify \$junos-underlying-interface as the predefined variable to represent the name of the underlying Ethernet interface on which the router creates the dynamic PPPoE logical interface. The variable is dynamically replaced with the actual name of the underlying interface supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

6. Configure the router to act as a PPPoE server when a PPPoE logical interface is dynamically created.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

7. Configure the IPv4 family for the pp0 interface. Specify the unnumbered address to dynamically create loopback interfaces. Because the example uses routing instances, assign the predefined variable `$junos-loopback-interface`.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet unnumbered-address $junos-loopback-interface
```

8. Configure the IPv6 family for the pp0 interface. Specify the unnumbered address to dynamically create loopback interfaces. Because the example uses routing instances without router advertisement, assign the predefined variable `$junos-loopback-interface`.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 unnumbered-address $junos-loopback-interface
```

9. Configure one or more PPP authentication protocols for the pp0 interface.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
user@host# set ppp-options pap
```

10. Enable keepalives and set an interval for keepalives. We recommend an interval of 30 seconds.

```
[edit dynamic-profiles pppoe-subscriber-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set keepalives interval 30
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit dynamic-profiles pppoe-subscriber-profile]
user@host# show
routing-instances {
  "$junos-routing-instance" {
    interface "$junos-interface-name";
  }
}
interfaces {
```

```

pp0 {
    unit "$junos-interface-unit" {
        ppp-options {
            chap;
            pap;
        }
        pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
        }
        keepalives interval 30;
        family inet {
            unnumbered-address "$junos-loopback-interface";
        }
        family inet6 {
            unnumbered-address "$junos-loopback-interface";
        }
    }
}

```

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

Configuring a Loopback Interface

CLI Quick Configuration

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

```

edit interfaces lo0
set unit 0 family inet address 203.0.113.1/32 primary
set unit 0 family inet address 203.0.113.1/32 preferred
set unit 0 family inet6 address 2001:db8:0::1/128 primary
set unit 0 family inet6 address 2001:db8:0::1/128 preferred

```

Step-by-Step Procedure

To configure a loopback interface:

1. Create the loopback interface and specify a unit number.

```
[edit]
user@host# edit interfaces lo0 unit 0
```

2. Configure the interface for IPv4.

```
[edit interfaces lo0 unit 0]
user@host# set family inet address 203.0.113.1/32 primary preferred
```

3. Configure the interface for IPv6.

```
[edit interfaces lo0 unit 0]
user@host# set family inet6 address 2001:db8:0::1/128 primary preferred
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces lo0]
user@host# show
unit 0 {
  family inet {
    address 203.0.113.1/32 {
      primary;
      preferred;
    }
  }
  family inet6 {
    address 2001:db8:0::1/128 {
      primary;
      preferred;
    }
  }
}
```

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

Configuring a VLAN Demux Interface over an Ethernet Underlying Interface

CLI Quick Configuration

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

```
edit interfaces
set ge-0/3/0 hierarchical-scheduler maximum-hierarchy-levels 2
set ge-0/3/0 flexible-vlan-tagging
set ge-0/3/0 encapsulation flexible-ethernet-services
exit
edit interfaces demux0 unit 1
set vlan-tags outer 1
set vlan-tags inner 1
set demux-options underlying-interface ge-0/3/0
set family pppoe dynamic-profile pppoe-subscriber-profile
set family pppoe duplicate-protection
set proxy-arp
```

Step-by-Step Procedure

To configure a VLAN demux interface over an Ethernet underlying interface:

1. Configure the underlying Ethernet interface.

```
[edit]
user@host# edit interfaces ge-0/3/0
user@host# set flexible-vlan-tagging
user@host# set encapsulation flexible-ethernet-services
user@host# set hierarchical-scheduler maximum-hierarchy-levels 2
```

2. Create the VLAN demux interface, and specify a unit number.

```
[edit]
user@host# edit interfaces demux0 unit 1
```

3. Configure the VLAN tags.

```
[edit interfaces demux0 unit 1]
user@host# set vlan-tags outer 1 inner 1
```

4. Specify the underlying Ethernet interface.

```
[edit interfaces demux0 unit 1]
user@host# set demux-options underlying-interface ge-0/3/0
```

5. Specify the dynamic profile.

```
[edit interfaces demux0 unit 1]
user@host# set family pppoe dynamic-profile pppoe-subscriber-profile
```

6. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on the same VLAN interface.

```
[edit interfaces demux0 unit 1]
user@host# set family pppoe duplicate-protection
```

7. (Optional) Specify that you want the demux interface to use Proxy ARP.

```
[edit interfaces demux0 unit 1]
user@host# set proxy-arp
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces]
user@host# show
ge-0/3/0 {
    hierarchical-scheduler maximum-hierarchy-levels 2;
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
}
```

```

demux0 {
    unit 1 {
        proxy-arp;
        vlan-tags outer 1 inner 1;
        demux-options {
            underlying-interface ge-0/3/0;
        }
        family pppoe {
            duplicate-protection;
            dynamic-profile pppoe-subscriber-profile;
        }
    }
}

```

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

Configuring an Interface for Communication with RADIUS Server

CLI Quick Configuration

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

```

edit interfaces ge-0/2/5
set unit 0 family inet address 203.0.113.99
set gigether-options no-auto-negotiation

```

Step-by-Step Procedure

To configure the interface:

1. Create the interface, specify a unit number, and configure the address.

```

[edit]
user@host# edit interfaces ge-0/2/5

```


2. Configure the interface for IPv4 and specify the address.

```
[edit interfaces ge-0/2/5]
user@host# set unit 0 family inet address 203.0.113.99
```

3. Specify that Gigabit Ethernet options are not automatically negotiated.

```
[edit interfaces ge-0/2/5]
user@host# set gigether-options no-auto-negotiation
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces ge-0/2/5]
user@host# show
gigether-options {
    no-auto-negotiation;
}
unit 0 {
    family inet {
        address 203.0.113.99/32;
    }
}
```

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

Specifying the BNG IP Address

CLI Quick Configuration

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

```
edit routing-options
set router-id 203.0.113.0
```



BEST PRACTICE: We strongly recommend that you configure the BNG IP address, thereby avoiding unpredictable behavior if the interface address on a loopback interface changes.

Step-by-Step Procedure

To configure the IP address of the BNG:

1. Access the routing-options configuration.

```
[edit]  
user@host# edit routing-options
```

2. Specify the IP address or the BNG.

```
[edit routing-options]  
user@host# set router-id 203.0.113.0
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit routing-options]  
user@host# show  
router-id 203.0.113.0;
```

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

Configuring RADIUS Server Access

CLI Quick Configuration

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

```
edit access radius-server 203.0.113.99
set secret "$ABC123$ABC123$ABC123"
set timeout 45
set retry 4
set source-address 203.0.113.1
```

Step-by-Step Procedure

To configure RADIUS servers:

1. Create a RADIUS server configuration, and specify the address of the server.

```
[edit]
user@host# edit access radius-server 203.0.113.99
```

2. Configure the required secret (password) for the server. Secrets enclosed in quotation marks can contain spaces.

```
[edit access radius-server 203.0.113.99]
user@host# set secret "$ABC123$ABC123$ABC123"
```

3. Configure the source address that the BNG uses when it sends RADIUS requests to the RADIUS server.

```
[edit access radius-server 203.0.113.99]
user@host# set source address 203.0.113.1
```

4. (Optional) Configure the number of times that the router attempts to contact a RADIUS accounting server. You can configure the router to retry from 1 through 16 times. The default setting is 3 retry attempts.

```
[edit access radius-server 203.0.113.99]
user@host# set retry 4
```

5. (Optional) Configure the length of time that the local router or switch waits to receive a response from a RADIUS server. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

```
[edit access radius-server 203.0.113.99]
user@host# set timeout 45
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
radius-server {
  203.0.113.99 {
    secret "$ABC123$ABC123$ABC123"; ## SECRET-DATA
    timeout 45;
    retry 4;
    source-address 203.0.113.1;
  }
}
```

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

Configuring RADIUS Server Access Profile

CLI Quick Configuration

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

```
edit access profile Access-Profile
set authentication-order radius
set radius authentication-server 203.0.113.99
set radius accounting-server 203.0.113.99
set accounting order radius
set accounting order none
set accounting update-interval 120
set accounting statistics volume-time
```

Step-by-Step Procedure

To configure a RADIUS server access profile:

1. Create a RADIUS server access profile.

```
[edit]
user@host# edit access profile Access-Profile
```

2. Specify the order in which authentication methods are used.

```
[edit access profile Access-Profile]
user@host# set authentication-order radius
```

3. Specify the address of the RADIUS server used for authentication and the server used for accounting.

```
[edit access profile Access-Profile]
user@host# set radius authentication-server 203.0.113.99
user@host# set radius accounting-server 203.0.113.99
```

4. Configure RADIUS accounting values for the access profile.

```
[edit access profile Access-Profile]
user@host# set accounting order [ radius none ]
user@host# set accounting update-interval 120
user@host# set accounting statistics volume-time
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
profile Access-Profile {
  authentication-order radius;
  radius {
    authentication-server 203.0.113.99;
    accounting-server 203.0.113.99;
  }
  accounting {
    order [ radius none ];
    update-interval 120;
    statistics volume-time;
  }
}
```

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

Configuring Local Address-Assignment Pools

CLI Quick Configuration

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

```
edit access address-assignment
set pool v4-pool family inet network 203.0.113.161/32
set pool v4-pool family inet range v4-range-0 low 203.0.113.161
```

```

set pool v4-pool family inet range v4-range-0 high 203.0.113.255
set pool v4-pool family inet dhcp-attributes maximum-lease-time 99999
set pool v6-ia-na-pool family inet6 prefix 2001:db8:1000:0000::/64
set pool v6-ia-na-pool family inet6 range v6-range-0 low 2001:db8:1000::1/128
set pool v6-ia-na-pool family inet6 range v6-range-0 high 2001:db8:1000::ffff:ffff/128
set pool v6-pd-pool family inet6 prefix 2001:db8:2012::/48
set pool v6-pd-pool family inet6 range v6-pd prefix-length 64

```

Step-by-Step Procedure

Configure three address-assignment pools for DHCPv4, DHCPv6 IA_NA, and DHCPv6 prefix delegation.

To configure the address-assignment pools:

1. Configure the address-assignment pool for DHCPv4.

```

[edit]
user@host# edit access address-assignment pool v4-pool
user@host# edit family inet
user@host# set network 203.0.113.161
user@host# set range v4-range-0 low 203.0.113.161
user@host# set range v4-range-0 high 203.0.113.255
user@host# set dhcp-attributes maximum-lease-time 99999

```

2. Configure the address-assignment pool for DHCPv6 IA_NA.

```

[edit]
user@host# edit access address-assignment pool v6-ia-na-pool
user@host# edit family inet6
user@host# set prefix 2001:db8:1000:0000::/64
user@host# set range v6-range-0 low 2001:db8:1000::1/128
user@host# set range v6-range-0 high 2001:db8:1000::ffff:ffff/128

```

3. Configure the address-assignment pool for DHCPv6 prefix delegation.

```

[edit]
user@host# edit access address-assignment pool v6-pd-pool
user@host# edit family inet6

```

```
user@host# set prefix 2001:db8:2012::/48
user@host# set range v6-pd prefix-length 64
```

4. (Optional) Enable duplicate prefix protection.

```
[edit access]
user@host# set address-protection
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
  address-assignment {
    pool v4-pool {
      family inet {
        network 203.0.113.161/32;
        range v4-range-0 {
          low 203.0.113.161;
          high 203.0.113.255;
        }
        dhcp-attributes {
          maximum-lease-time 99999;
        }
      }
    }
    pool v6-ia-na-pool {
      family inet6 {
        prefix prefix 2001:db8:1000:0000::/64 ;
        range v6-range-0 {
          low 2001:db8:1000::1/128;
          high 2001:db8:1000::ffff:ffff/128;
        }
      }
    }
    pool v6-pd-pool {
      family inet6 {
        prefix 2001:db8:2012::/48;
        range v6-pd prefix-length 64;
      }
    }
  }
```



```

    }
  }
}

address-protection;
}

```

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

Verification

IN THIS SECTION

- [Verifying Active Subscriber Sessions | 1298](#)
- [Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | 1299](#)
- [Verifying Dynamic Subscriber Sessions | 1300](#)
- [Verifying DHCPv6 Address Pools Used for DHCPv6 Prefix Delegation | 1301](#)
- [Verifying DHCPv6 Address Bindings | 1302](#)
- [Verifying PPP Options Negotiated with the Remote Peer | 1303](#)

Confirm that the configuration is working properly.

Verifying Active Subscriber Sessions

Purpose

Verify active subscriber sessions.

Action

From operational mode, enter the `show subscribers summary` command.

```

user@host>show subscribers summary
Subscribers by State
  Active: 2
  Total: 2

```

Subscribers by Client Type	
DHCP:	1
PPPoE:	1
Total:	2

Meaning

The fields under Subscribers by State show the number of active subscribers.

The fields under Subscribers by Client Type show the number of active DHCP and PPPoE subscriber sessions.

Verifying Both IPv4 and IPv6 Address in Correct Routing Instance

Purpose

Verify that the subscriber has both an IPv4 and IPv6 address and is placed in the correct routing instance.

Action

From operational mode, enter the show subscribers command.

user@host>show subscribers			
Interface	IP Address/VLAN ID	User Name	LS:RI
pp0.1073741825	203.0.113.162	SBRSTATICUSER	default:default
pp0.1073741825	2001:db8:1000::1		default:default

Meaning

The Interface field shows that two subscriber sessions are running on the same interface. The IP Address field shows that one session is assigned an IPv4 address, and the second session is assigned an IPv6 address by DHCPv6 IA_NA.

The LS:RI field shows that the subscriber is placed in the correct routing instance and that traffic can be sent and received.

Verifying Dynamic Subscriber Sessions

Purpose

Verify dynamic PPPoE and DHCPv6 subscriber sessions. In this example configuration the DHCPv6 subscriber session should be layered over the underlying PPPoE subscriber session.

Action

From operational mode, enter the `show subscribers detail` command.

```
user@host>show subscribers detail
Type: PPPoE
User Name: SBRSTATICUSER
IP Address: 203.0.113.162
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: pppoe-subscriber-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-12-08 09:11:41 PST

Type: DHCP
IPv6 Address: 2001:db8:1000::1
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5E:00:53:31
State: Active
Radius Accounting ID: 3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-12-08 09:12:11 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 01 02 00 00 01
```

```
00 06 00 02 00 03 00 03 00 0c 00 00 00 00 00 00 00 00 00
00 00
```

Meaning

When a subscriber has logged in and started both an IPv4 and an IPv6 session, the output shows the active underlying PPPoE session and the active DHCPv6 session.

The Session ID field for the PPPoE session is 2. The Underlying Session ID for the DHCP session is 2, which shows that the PPPoE session is the underlying session.

Verifying DHCPv6 Address Pools Used for DHCPv6 Prefix Delegation

Purpose

Verify the delegated address pool used for DHCPv6 prefix delegation and the length of the IPv6 prefix that was delegated to the CPE.

Action

From operational mode, enter the `show subscribers extensive` command.

```
user@host>show subscribers extensive
Type: PPPoE
User Name: SBRSTATICUSER
IP Address: 203.0.113.162
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: pppoe-subscriber-profile
MAC Address: 00:00:5E:00:53:31
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-12-08 09:11:41 PST
IPv6 Delegated Address Pool: v6-na-pool

Type: DHCP
IPv6 Address: 2001:db8:1000::1
```

```

Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5E:00:53:31
State: Active
Radius Accounting ID: 3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-12-08 09:12:11 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 01 02 00 00 01
00 06 00 02 00 03 00 03 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Address Pool: v6-na-pool
IPv6 Delegated Network Prefix Length: 64

```

Meaning

The IPv6 Delegated Address Pool field shows the name of the pool that DHCPv6 used to assign the IPv6 address for this subscriber session.

Verifying DHCPv6 Address Bindings

Purpose

Display the address bindings in the client table on the DHCPv6 local server.

Action

From operational mode, enter the `show dhcpv6 server binding detail` command.

```

user@host>show dhcpv6 server binding detail
Session Id: 580547
    Client IPv6 Address:                2001:db8:1000::4/128
    Client DUID:                        LL0x1-00:01:02:00:00:01
    State:                              BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUN
D)
    Lease Expires:                      2012-01-05 07:06:04 PST
    Lease Expires in:                   82943 seconds
    Lease Start:                        2012-01-04 07:06:04 PST

```

```

Last Packet Received:      2012-01-04 07:06:04 PST
Incoming Client Interface:  pp0.1073926645
Server Ip Address:         0.0.0.0
Client Pool Name:          v6-na-pool-0
Client Id Length:          10
Client Id:                 /0x00030001/0x00010200/0x0001

```

Meaning

The **Client IPv6 Address** field shows the /128 address that was assigned to the CPE WAN link using DHCPv6 IA_NA.

The **Client Pool Name** field shows the name of the address pool that was used to assign the **Client IPv6 Address**.

Verifying PPP Options Negotiated with the Remote Peer

Purpose

Verify PPP options negotiated with the remote peer.

Action

From operational mode, enter the show ppp interface *interface* extensive command.

```

user@host>show ppp interface pp0.1073741825 extensive
Session pp0.1073926645, Type: PPP, Phase: Network
LCP
  State: Opened
  Last started: 2012-01-04 07:05:33 PST
  Last completed: 2012-01-04 07:05:33 PST
  Negotiated options:
    Authentication protocol: pap, Magic number: 191301485, Local MRU: 1492,
    Peer MRU: 65531
Authentication: PAP
  State: Grant
  Last started: 2012-01-04 07:05:33 PST
  Last completed: 2012-01-04 07:05:33 PST
IPCP
  State: Opened

```

```

Last started: 2012-01-04 07:05:34 PST
Last completed: 2012-01-04 07:05:34 PST
Negotiated options:
  Local address: 203.0.113.1, Remote address: 203.0.113.162

```

```

IPV6CP
State: Opened
Last started: 2012-01-04 07:05:34 PST
Last completed: 2012-01-04 07:05:34 PST
Negotiated options:
  Local interface identifier: 2a0:a50f:fc71:e049,
  Remote interface identifier: 201:2ff:fe00:1

```

Meaning

The output shows the PPP options that were negotiated with the remote peer.

Under IPCP, the Negotiated options field shows the IPv4 local and remote addresses that were negotiated by IPCP.

Under IPV6CP, the Negotiated options field shows the IPv6 local and remote interface identifier that were negotiated by IPV6CP.

SEE ALSO

[Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation](#) | 1308

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.3R1	Starting in Junos OS Release 17.2R3, 17.4R2, 18.1R3, 18.2R2, and 18.3R1, the jdhcpd process extends the lease for both address types in this situation.

RELATED DOCUMENTATION

[IPv6 WAN Link Addressing with DHCPv6 IA_NA](#) | 1257

[Subscriber LAN Addressing with DHCPv6 Prefix Delegation](#) | 1259

Designs for IPv6 Addressing in a Subscriber Access Network

IN THIS SECTION

- [Selecting the Type of Addressing Used on the CPE | 1305](#)
- [Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link | 1306](#)
- [Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | 1306](#)
- [Selecting the Method of Obtaining IPv6 Prefixes | 1307](#)
- [Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 1308](#)
- [Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | 1309](#)
- [Design 3: IPv6 Addressing with NDRA | 1310](#)
- [Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | 1311](#)

Selecting the Type of Addressing Used on the CPE

In some networks, you do not need to assign a global IPv6 address on the CPE WAN link. Your decision depends on the type of CPE being used:

- If the CPE is purchased by the subscriber, and is not a device specifically recommended by the service provider, you need to assign a global IPv6 address that can be routed on the Internet.
- If the CPE is supplied by or recommended by the service provider, you can use the loopback interface to manage the CPE.

In this case, you can use a link-local address or you can use an address that is derived from DHCPv6 prefix delegation.

SEE ALSO

| [IPv6 Addressing Requirements for a Subscriber Access Network | 1248](#)

Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link

To assign a global IPv6 address to the WAN link of the CPE device, you can choose one of the methods described in [Table 56 on page 1306](#).

Table 56: Choosing the Global IPv6 Address Provisioning Method for the WAN Link

NDRA Features	DHCPv6 IA_NA Features
Provides address autoconfiguration of the WAN link by means of router advertisements.	Provides a single IPv6/128 address to the WAN interface of the CPE by the BNG acting as a DHCPv6 server.
Supported on PPPoE access networks.	Supported on PPPoE and DHCP access networks.
Provides duplicate prefix prevention.	Provides the ability to use one DHCPv6 message to solicit both a global IPv6 address for the WAN link, and a prefix used to provision addresses on the subscriber LAN.
Use if the CPE does not support DHCP.	--

SEE ALSO

[IPv6 Addressing Requirements for a Subscriber Access Network | 1248](#)

[Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250](#)

[IPv6 WAN Link Addressing with DHCPv6 IA_NA | 1257](#)

Selecting the Method of Assigning Global IPv6 Addresses to Subscribers



BEST PRACTICE: For addressing on the subscriber LAN, we recommend that you provision a global IP address for each device on the LAN. IPv6 was designed to allow every IP-capable device on a subscriber LAN to obtain a globally unique address, which avoids the use of NAT between the subscriber LAN and the service provider.

DHCPv6 prefix delegation automates the delegation of IPv6 prefixes to the CPE. The CPE can then use these prefixes to assign global IPv6 addresses for use in a subscriber LAN. DHCPv6 prefix delegation is useful when the delegating router (the BNG) does not have information about the topology of the networks in which the requesting router (the CPE) is located. In such cases, the delegating router requires only the identity of the requesting router to choose a prefix for delegation.

SEE ALSO

- Using DHCPv6 Prefix Delegation Overview | 1260
- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269

Selecting the Method of Obtaining IPv6 Prefixes

IN THIS SECTION

- Using a AAA RADIUS Server to Obtain Global IPv6 Addresses and IPv6 Prefixes | 1307
- Using a Local Pool to Assign IPv6 Addresses or Prefixes | 1308

You can set up the BNG to select IPv6 prefixes through one of the following methods:

- An external source such as a AAA RADIUS server or a DHCP server using the DHCPv6 relay agent.
- Dynamic assignment from a local pool of global IPv6 addresses or prefixes that is configured on the BNG

Using a AAA RADIUS Server to Obtain Global IPv6 Addresses and IPv6 Prefixes

Table 57 on page 1307 describes the RADIUS attributes used in a dual-stack network. These attributes are sent from the RADIUS server to the BNG in RADIUS Access-Accept messages.

Table 57: RADIUS Attributes Used to Obtain Global IPv6 Addresses and IPv6 Prefixes

RADIUS Attribute	Address Assignment Type	Attribute Description
Framed-IPv6-Prefix	NDRA	IPv6 prefix with a prefix length less than 128.
	DHCPv6 IA_NA	IPv6 prefix with a length of 128.
Framed-IPv6-Pool	NDRA	Name of an NDRA pool configured on the BNG from which the BNG selects a prefix.
	DHCPv6 IA_NA	Name of an address-assignment pool configured on the BNG from which the BNG selects a global IPv6 address.

Table 57: RADIUS Attributes Used to Obtain Global IPv6 Addresses and IPv6 Prefixes (Continued)

RADIUS Attribute	Address Assignment Type	Attribute Description
Delegated-IPv6-Prefix	DHCPv6 prefix delegation	IPv6 prefix.
IPv6-Delegated-Pool-Name	DHCPv6 prefix delegation	Name of an address-assignment pool configured on the BNG from which the BNG delegates a prefix.

Using a Local Pool to Assign IPv6 Addresses or Prefixes

You can use the `delegated-pool` statement to specify a local address-assignment pool on the BNG to provide delegated prefix addresses at any of the following hierarchy levels:

- Globally for the server at the `[edit system services dhcp-local-server dhcpv6 overrides]` hierarchy level.
- For a named group of interfaces at the `[edit system services dhcp-local-server dhcpv6 group group-name overrides]` hierarchy level.
- For a specific interface within a named group of interface at the `[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name overrides]` hierarchy level.



NOTE: A pool specified by the Juniper Networks IPv6-Delegated-Pool-Name VSA (26-161) takes precedence over a locally configured pool.

SEE ALSO

[IPv6 Addressing Requirements for a Subscriber Access Network | 1248](#)

[Using DHCPv6 Prefix Delegation Overview | 1260](#)


Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

This design ([Figure 13 on page 1309](#)) uses DHCPv6 IA_NA and DHCPv6 prefix delegation in your subscriber access network as follows:

- DHCPv6 IA_NA is used to assign a global IPv6 address on the WAN link. The address can come from a local pool or AAA RADIUS.

- DHCPv6 prefix delegation is used for host device addressing. The delegated prefix can come from a local pool or from AAA RADIUS. The CPE uses the delegated prefix for subscriber addressing. The CPE can use NDRA or DHCPv6 to allocate IPv6 addresses on the LAN.

Figure 13: Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation



SEE ALSO

[Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | 1274](#)

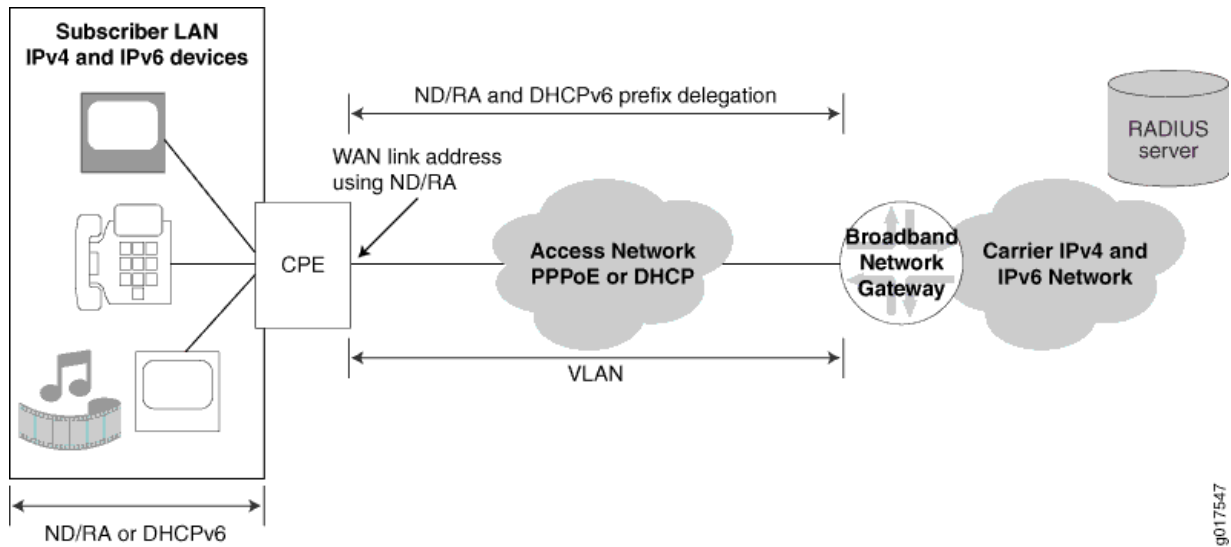
[Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269](#)

Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation

This design ([Figure 14 on page 1310](#)) uses NDRA and DHCPv6 prefix delegation in your subscriber access network as follows:

- NDRA addressing is used to provision a global IPv6 address on the WAN link. IPv6 prefixes for NDRA come from a local pool or AAA RADIUS.
- DHCPv6 prefix delegation is used for host device addressing. The delegated prefix can come from a local pool or from AAA RADIUS. The CPE uses the delegated prefix for subscriber addressing. The CPE can use NDRA or DHCPv6 to allocate IPv6 addresses on the LAN.

Figure 14: Subscriber Access Network with NDRA and DHCPv6 Prefix Delegation



If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of design 2 and design 3.

SEE ALSO

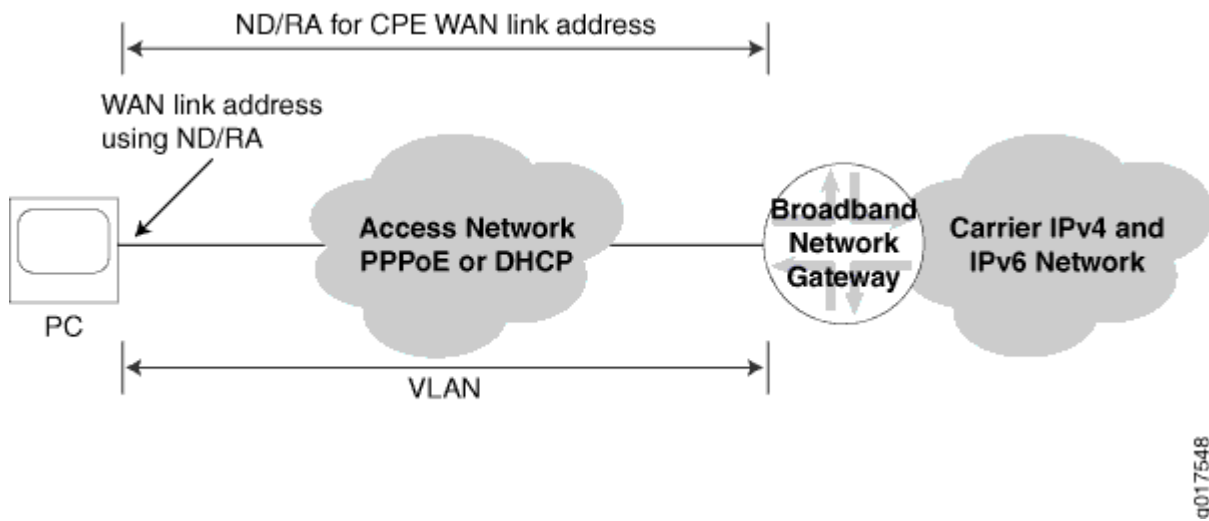
[Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | 1390](#)

[Using DHCPv6 Prefix Delegation Overview | 1260](#)

Design 3: IPv6 Addressing with NDRA

In this design ([Figure 15 on page 1311](#)), NDRA is used for addressing a global IPv6 on the WAN link with prefixes from a local pool or AAA RADIUS. The PC does not need a delegated prefix.

Figure 15: Subscriber Access Network with NDRA



If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of Design 2 and Design 3.


SEE ALSO

[Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | 1364](#)

Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix

In this design ([Figure 16 on page 1312](#)), the CPE is a model that is sold by or specified by the service provider. The CPE uses an unnumbered WAN interface. The BNG delegates an IPv6 prefix to the CPE with DHCPv6 prefix delegation. The CPE uses the delegated prefix for subscriber addressing. It can use NDRA or DHCPv6 to allocate the IPv6 addresses on the LAN.

Figure 16: Subscriber Access Network with DHCPv6 Prefix Delegation

**SEE ALSO**

[Using DHCPv6 Prefix Delegation Overview](#) | 1260

RELATED DOCUMENTATION

[Migration to IPv6 Using IPv4 and IPv6 Dual Stack](#) | 1245

[Subscriber LAN Addressing with DHCPv6 Prefix Delegation](#) | 1259

[IPv6 WAN Link Addressing with NDRA](#) | 1250

[Dual Stack for PPPoE Access Networks Using NDRA](#) | 1360

[WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation](#) | 1269

Dual-Stack Access Models in a DHCP Network**IN THIS SECTION**

● [IPv4 and IPv6 Dual Stack in a DHCP Access Network](#) | 1313

- AAA Service Framework in a Dual Stack over a DHCP Access Network | 1314
- Dual-Stack Interface Stack in a DHCP Wholesale Network | 1315
- Single-Session DHCP Dual-Stack Overview | 1316
- Configuring Single-Session DHCP Dual-Stack Support | 1320
- Verifying and Managing DHCP Dual-Stack Configuration | 1323

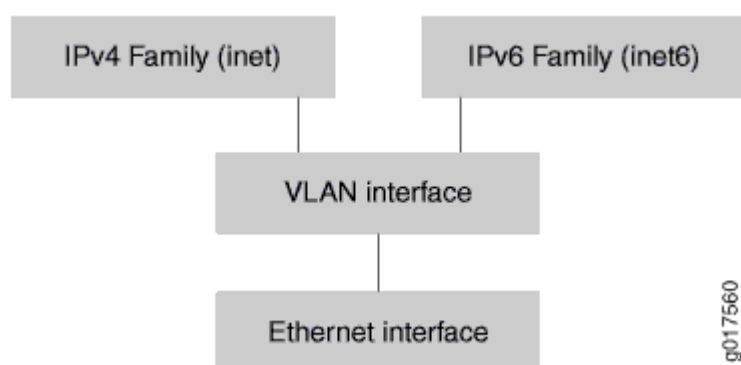
IPv4 and IPv6 Dual Stack in a DHCP Access Network

IN THIS SECTION

- Support for Demultiplexing Interfaces | 1314

Figure 17 on page 1313 shows a dual-stack interface stack in a DHCP access network. The IPv4 family (inet) and the IPv6 family (inet6) can reside on the same VLAN interface.

Figure 17: Dual-Stack Interface Stack over a DHCP Access Network



NOTE: When you are using IPv4 and IPv6 dual stack on the same DHCP interface, you must configure one dynamic profile for both the IPv4 and IPv6 subscribers. You cannot run IPv4 and IPv6 subscriber sessions over the same interface if you configure separate dynamic profiles for IPv4 and IPv6.

Support for Demultiplexing Interfaces

IPv4 and IPv6 dual stack is supported on VLAN demultiplexing (demux) interfaces. Dual stack is not supported on IP demux interfaces.

SEE ALSO

| [Basic Architecture of a Subscriber Access Dual-Stack Network](#) | 1245

AAA Service Framework in a Dual Stack over a DHCP Access Network

IN THIS SECTION

- [Collection of Accounting Statistics in a DHCP Access Network](#) | 1315
- [Change of Authorization \(CoA\)](#) | 1315

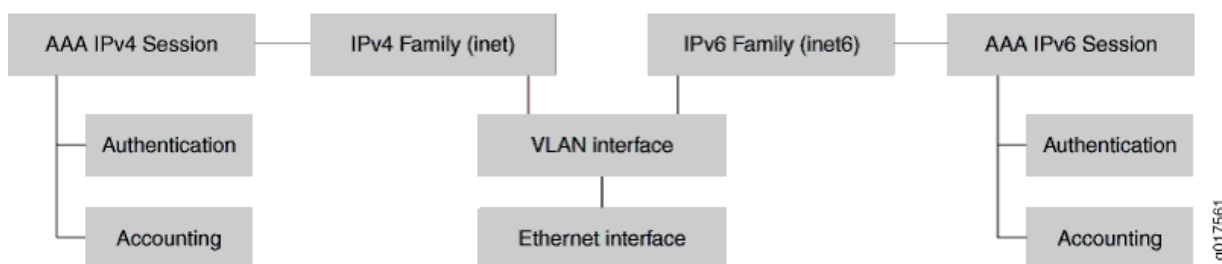
You can use the AAA Service Framework for all authentication, authorization, accounting, address assignment, and dynamic request services that the BNG uses for network access. The framework supports authentication and authorization through external RADIUS servers. It also supports accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS servers.

The BNG interacts with external servers to determine how individual subscribers access the broadband network. The BNG can also obtain information from external servers for the following:

- How subscribers are authenticated.
- How accounting statistics are collected and used.
- How dynamic requests, such as CoA, are handled.

As shown in [Figure 18 on page 1315](#), an implementation of dual stack over a DHCP access network, there are separate AAA sessions for IPv4 and IPv6 authentication and accounting.

Figure 18: AAA Service Framework in a Dual Stack over a DHCP Access Network



Collection of Accounting Statistics in a DHCP Access Network

AAA provides support for IPv4 and IPv6 statistics in separate accounting sessions.

The following RADIUS attributes are included by default (when available) in Acct-Start, Interim, and Acct-Stop messages:

- Framed-IPv6-Prefix
- Framed-IPv6-Pool
- Delegated-Ipv6-Prefix
- Framed-IPv4-Route
- Framed-IPv6-Route

You can configure the BNG to exclude these attributes in accounting Acct-Start and Acct-Stop messages.

Change of Authorization (CoA)


RADIUS servers can initiate dynamic requests to the BNG. Dynamic requests include CoA requests, which specify VSA modifications and service changes.

In your access profile configuration, you specify the IP addresses of RADIUS authentication servers that can initiate dynamic requests to the router. The list of authentication servers also provides RADIUS-based dynamic service activation and deactivation during subscriber login.

Dual-Stack Interface Stack in a DHCP Wholesale Network

[Figure 19 on page 1316](#) shows a dual-stack interface stack in a DHCP wholesale network. In this scenario, the IPv4 and IPv6 demux interfaces are configured on the same VLAN interface. The demux interfaces are configured in a separate logical system: routing instance.

Figure 19: Dual-Stack Interface Stack in a DHCP Wholesale Network



Single-Session DHCP Dual-Stack Overview

Junos OS supports a single-session DHCP dual-stack, which simplifies management of dual-stack subscribers, and improves performance and session requirements when compared to the traditional dual-stack support.


In a DHCP dual-stack environment, a DHCP server supports both DHCPv4 and DHCPv6 subscribers. The DHCP server provides services, such as authentication and accounting, for both the DHCPv4 and DHCPv6 legs of the dual-stack. In a traditional implementation, the two legs of the dual-stack are viewed as being independent. The presence of separate legs for DHCPv4 and DHCPv6 creates inefficiencies, since separate, and multiple, sessions can be required to provide similar support for each leg of the dual-stack. For example, to provide authentication for a traditional dual-stack over a dynamic VLAN requires three separate sessions, one for DHCPv4, one for DHCPv6, and one for the authenticated dynamic VLAN. Similarly, multiple sessions might be also required for dual-stack accounting operations.

In the dual-stack over a dynamic VLAN, the single-session dual-stack requires only a single session for authentication, as opposed to the three sessions required for the traditional dual-stack configuration. Accounting support for the dual-stack also uses a single session. In addition to reducing the number of sessions required, the single-session feature also simplifies router configuration, reduces RADIUS message load, and improves accounting session performance for households with dual-stack environments.

In the single-session dual-stack environment, the first DHCP session that negotiates will trigger the dynamic VLAN creation (if required) and is authorized at the DHCP application. The second leg of the dual-stack is held off until the authorization point is complete. When the second leg of the dual stack is established, the DHCP client inherits all common subscriber database values, such as circuit-id, remote-id, username, and interface name from the first leg.

In [Figure 20 on page 1317](#), single subscriber session is established for dual-stack user.

Figure 20: DHCP Dual Stack Single-Session Subscriber Deployment Model



You can configure single-session dual-stack subscriber settings for DHCP relay agent and DHCP local server. You use the `dual-stack-group` statement to create a named group that specifies the values for dual stack subscribers. Then, you use the `dual-stack` statement to specify the name of the dual stack group and assign the group to subscribers at the global, group, or interface level.

- For DHCP relay agent, configure these statements at the `[edit forwarding-options dhcp-relay]` hierarchy level and the `[edit forwarding-options dhcp-relay ... overrides]` hierarchy level, respectively,
- For DHCP local server, configure these statements at the `[edit system services dhcp-local-server]` hierarchy level and the `[edit system services dhcp-local-server ... overrides]` hierarchy level, respectively,

You can configure the following common DHCP settings for the single-session dual-stack model. In most cases, these settings are similar to those used for separate DHCPv4 and DHCPv6 legs in a traditional dual-stack configuration. When configured and referenced, the dual-stack configuration takes precedence over the same items configured under the respective family.

- **access-profile**—Access profile that provides authentication and accounting parameters for the dual-stack group that take precedence over those configured in a global access profile or in a profile configured for the DHCP relay agent or DHCP local server.
- **authentication**—Authentication-related parameters (such as password and username) the router sends to the external AAA server.

The dual-stack authentication stanza is similar to the stanza available separately for the v4 and v6 address families. When the `username-include` configuration syntax is used for the DHCPv4 leg of the dual-stack, the `relay-agent-interface-id` option is equivalent to the DHCPv4 `relay-option-82 circuit-id` statement, and the `relay-agent-remote-id` option is equivalent to the DHCPv4 `relay-option-82 remote-id` statement. You do not have to configure the two DHCPv4 options separately.

- `classification-key`—Classification key defines mechanism to be used to identify a dual stack household.
- `dual-stack-interface-client-limit`—Limits the number of dual stack subscribers login per interface.



NOTE: For dual-stack subscribers, always use this statement instead of the `interface-client-limit` statement.

- `dynamic-profile`—Dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.
- `liveness-detection`—Configure an active liveness detection protocol that deletes the binding and releases the resources if the subscriber fails to respond to a configured number of consecutive liveness detection requests, the subscriber.
- `on-demand-address-allocation`—(DHCP local server) Designates whether on-demand address allocation mode is forced for a dual-stack subscriber.

If this configuration is not present, all IP addresses and prefixes for IPv4 and IPv6 families of a dual stack subscriber will be preallocated when the first leg of a dual stack subscriber initially logs in.

If this configuration is present when the first leg of a dual-stack subscriber initially logs in, RADIUS authentication is performed (if configured) and the IP address and prefix of this first family only will be allocated. The IP address and prefix for the other family will not be allocated unless the other family leg subsequently initially logs in.



NOTE: The IP address allocation for the second family is informed by the RADIUS authentication previously performed at the time of the first family login.

Starting in Junos OS Release 18.4R1, the method of address allocation is checked to determine subsequent behavior when authd notifies the DHCP process that an address pool is deleted or being drained. [Table 58 on page 1319](#) describes the behavior.

Table 58: Behavior When Address Pool is Deleted or Drained

Address Allocation Method	Address Pool is Drained	Address Pool is Deleted
On demand	Family with address in pool is logged out gracefully when a DHCP renew or rebind message is received.	Family with address in pool is logged out immediately.
Preallocated	Addresses for both families are deleted gracefully when a DHCP renew or rebind message is received.	Addresses for both families are deleted immediately.

- **protocol-master**—This term designates either an IPv4 or IPv6 family as the primary family for a dual stack subscriber. The secondary family client binding login-in will be rejected until a valid client binding is in place for the primary family.



CAUTION: If the secondary family binding is logged out for any reason, then only the secondary family binding will be torn down.
If the primary family binding is logged out for any reason, then the corresponding bindings for both the primary and secondary families will be torn down.

- **reauthenticate**—(DHCP local server) Configure reauthentication of the subscriber to initiate change characteristics such as service activations/deactivations and attribute modifications.
- **relay-agent-interface-id**—(DHCP relay agent) Includes Relay Agent Interface-ID (option 18) in DHCPv6 packets destined for the DHCPv6 server. You can configure numerous options to specify what is included in the circuit ID value.

For the DHCPv4 leg of the dual-stack, this statement includes the DHCPv4 relay-option-82 circuit-id in packets destined for the DHCPv4 server.

- **relay-agent-remote-id**—(DHCP relay agent) Includes Relay Agent Remote-ID (option 37) in DHCPv6 packets destined for a DHCPv6 server. You can configure numerous options to specify what is included in the remote ID value.

For the DHCPv4 leg of the dual-stack, this statement includes the DHCPv4 relay-option-82 remote-id in packets destined for the DHCPv4 server.

- **service-profile**—Dynamic profile for the default subscriber service (or the default DHCP client management service), which is activated when the subscriber (or client) logs in.

- short-cycle protection—Detect and lock out short-lived client sessions and clients that repeatedly fail session negotiation to reduce resource usage associated with connection and authentication processing in highly scaled networks.

SEE ALSO

[RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 727](#)

[Configuring RADIUS Reauthentication for DHCP Subscribers | 738](#)

Configuring Single-Session DHCP Dual-Stack Support

Configuring single-session dual-stack support is a two-step process. You first create the dual-stack group that specifies the configuration parameters that are shared between the DHCPv4 and DHCPv6 legs of the DHCP dual stack. Then, you attach the dual-stack group to DHCP subscriber interfaces by overriding the default DHCP configurations for the DHCPv4 and DHCPv6 subscribers. You must reference the dual-stack group for both legs of the dual stack. If you attach the group to one leg only, the router rejects the other leg. You can attach the dual-stack group globally, for a specified DHCP group of interfaces, or for a specific interface.

To configure single-session dual-stack group support.

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

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

2. Create and name the dual-stack group.

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

3. Attach an access profile to the dual-stack group to override the corresponding authentication and accounting properties configured in a global access profile or DHCP relay agent access profile.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
user@host# set access-profile profile-name
```

See ["Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces" on page 1030](#).

4. Configure the authentication username values and password for the dual-stack group.

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

See ["Specifying Authentication Support" on page 1161](#).

- Configure the unique username.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication]
user@host# set username-include <username-include-configuration>
```

See ["Creating Unique Usernames for DHCP Clients" on page 1162](#).

- Configure the password that authenticates the username to the external authentication service.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication]
user@host# set password password-string
```

See ["Example-Configuring DHCP with External Authentication Server" on page 1164](#).

5. Specify the dynamic profile associated with the dual-stack group.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
user@host# set dynamic-profile <dynamic-profile configuration>
```

See *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*.

6. Specify the service profile associated with the dual-stack group.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
user@host# set service-profile dynamic-profile-name
```

See *Defining Various Levels of Services for DHCP Subscribers*.

7. Specify the relay-agent-interface-id for the dual-stack group.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
user@host# set relay-agent-interface-id <relay-agent-interface-id configuration>
```

See ["Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets" on page 1458](#).



NOTE: For the DHCPv4 leg of the dual-stack, this step specifies the Option 82 Agent Circuit ID (suboption 1) for DHCPv4 clients. See ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).

8. Specify the relay-agent-remote-id for the dual-stack group.

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
user@host# set relay-agent-remote-id <relay-agent-remote-id-configuration>
```

See ["Inserting DHCPv6 Remote-ID Option \(Option 37\) In DHCPv6 Packets" on page 1460](#).



NOTE: For the DHCPv4 leg of the dual-stack, this step specifies the Option 82 Agent Remote ID (suboption 2) for DHCPv4 clients. See ["Using DHCP Relay Agent Option 82 Information" on page 1078](#).

9. Use the override feature to override the default DHCP relay behavior and assign the dual-stack group to DHCPv4 and DHCPv6 clients. You must perform separate steps for each leg of the dual stack.
 - To assign the dual-stack group to DHCPv4 clients:

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

See ["Overriding the Default DHCP Relay Configuration Settings" on page 1036](#).

- To assign the dual-stack group to DHCPv6 clients:

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

See ["Overriding the Default DHCP Relay Configuration Settings" on page 1036](#).

10. (Optional) Verify your dual-stack group configuration for DHCPv4 and DHCPv6.

See ["Verifying and Managing DHCP Dual-Stack Configuration" on page 1323](#).

```
user@host> show dhcp relay binding
user@host> show dhcpv6 relay binding
user@host> show subscribers
```

Verifying and Managing DHCP Dual-Stack Configuration

IN THIS SECTION

- Purpose | 1323
- Action | 1323

Purpose

Display information related to the DHCP single-session dual-stack configuration.

Action

- To display DHCP relay agent binding information for dual-stack clients:

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

- To display DHCPv6 relay agent binding information for dual-stack clients:

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

- To display assigned IP4 and IPv6 addresses for DHCP dual-stack clients:

```
user@host>show subscribers
```

- To show IPv4 and IPv6 addresses for a specific session:

```
user@host>show network-access aaa subscribers session-id session-id session-id detail
```

- To all clear DHCPv4 relay bindings and associated DHCPv6 bindings for the dual-stack in the default routing instance. This command does not effect DHCPv6-only stacks that are not associated with the dual-stack.

```
user@host>clear dhcp relay binding dual-stack all
```

Alternatively, you can limit clearing to an address, VLAN interface, logical system, or routing instance.

- To clear all DHCPv6 relay bindings and associated DHCPv4 bindings for the dual-stack in the default routing instance. This command does not effect DHCPv4-only stacks that are not associated with the dual-stack.

```
user@host>clear dhcpv6 relay binding dual-stack all
```

Alternatively, you can limit clearing to an address, VLAN interface, logical system, or routing instance.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, the method of address allocation is checked to determine subsequent behavior when authd notifies the DHCP process that an address pool is deleted or being drained.

RELATED DOCUMENTATION

- [Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245](#)
- [Dual-Stack Access Models in a DHCP Network | 1312](#)

Dual-Stack Access Models in a PPPoE Network

IN THIS SECTION

- [IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1325](#)
- [Shared IPv4 and IPv6 Service Sessions on PPP Access Networks | 1328](#)
- [AAA Service Framework in a Dual Stack over a PPPoE Access Network | 1328](#)
- [RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers | 1331](#)
- [Accounting Messages for PPPoE Using NDRA Prefixes | 1332](#)
- [Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes | 1339](#)

- [Suppressing Accounting Information That Comes from AAA | 1349](#)
- [Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address | 1350](#)

IPv4 and IPv6 Dual Stack in a PPPoE Access Network

IN THIS SECTION

- [Support for Demultiplexing Interfaces | 1327](#)
- [Determining the Status of CPE in a PPPoE Access Network | 1327](#)
- [IPv6 Address Provisioning in the PPPoE Access Network | 1327](#)
- [Authentication in a PPPoE Access Network | 1327](#)
- [Negotiation of Network Control Protocols When Authorized Addresses Are Unavailable | 1327](#)

In a dual-stack architecture with a PPPoE access network that connects the CPE to the BNG, IPv4 and IPv6 connectivity are provided over a single PPP logical link. The PPP IPv4 control protocol (IPCP) and the IPv6 control protocol (IPv6CP) provide independent IPv4 and IPv6 connectivity over the logical link.

The BNG and the CPE handle both IPCP and IPv6CP identically and simultaneously over a single PPP connection. The BNG or the CPE can open and close any Network Control Protocol (NCP) session without affecting the other sessions. This capability allows for a dynamic setup where IPv4 (family inet) and IPv6 (family inet6) sessions can be brought up and down individually. As long as one family is active, the subscriber remains active.

[Figure 21 on page 1326](#) shows a dual-stack interface stack in a PPPoE access network. The IPv4 family (inet) and the IPv6 family (inet6) can reside on the same PPPoE logical interfaces. The family inet and family inet6 parts of dynamic profiles are applied, and services are activated when each individual family is negotiated.

Figure 21: Dual-Stack Interface Stack over a PPPoE Access Network


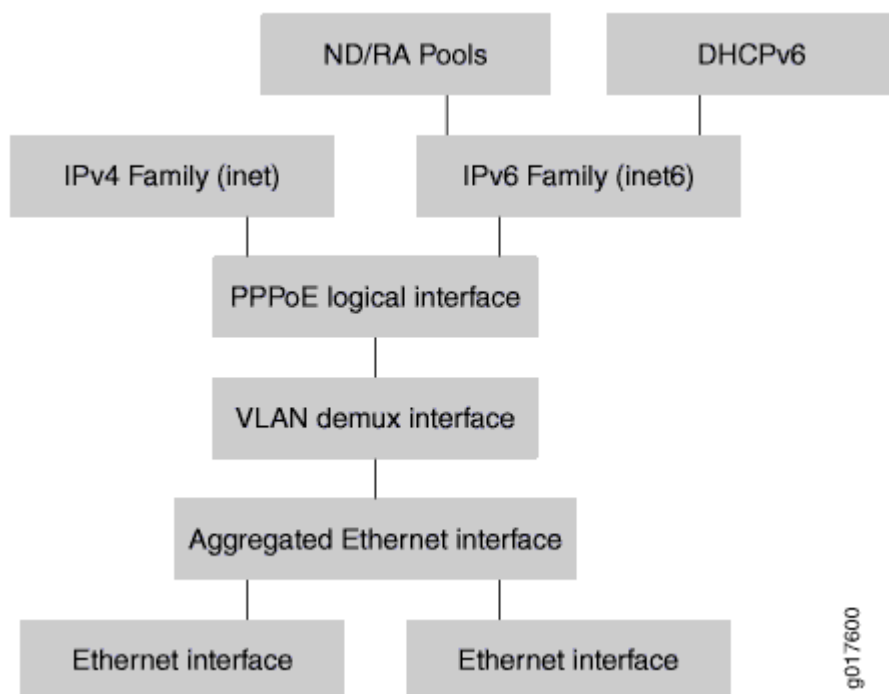


Figure 22 on page 1326 shows a dual-stack interface stack over aggregated Ethernet in a PPPoE access network.

Figure 22: Dual-Stack Aggregated Ethernet Stack over a PPPoE Access Network



Support for Demultiplexing Interfaces

IPv4 and IPv6 dual stack is supported on VLAN demultiplexing (demux) interfaces. Dual stack is not supported on IP demux interfaces.

Determining the Status of CPE in a PPPoE Access Network

In a PPPoE access network, you can enable keepalives to determine the status of the CPE.

IPv6 Address Provisioning in the PPPoE Access Network

IPv6CP negotiates the interface identifier, which can be used to provision link-local addresses that are used for direct connectivity between the BNG and the CPE. Because PPPoE negotiates only interface IDs and does not negotiate IPv6 addresses, PPPoE relies on other protocols for addressing. The protocols you can use are DHCPv6 and NDRA.

Authentication in a PPPoE Access Network

In a PPPoE network, you can use PAP and CHAP to identify and authenticate the CPE and subscriber sessions.

You can also use AAA for authentication and authorization through external RADIUS servers.

Negotiation of Network Control Protocols When Authorized Addresses Are Unavailable

NCP negotiation is initiated for subscriber sessions by default, even when authorized addresses are not available. An example of this situation is when the DHCPv6 local server is configured with an override so that the jpppd process never receives an IPv6 address or prefix from AAA, although the DHCPv6 local server receives a prefix from a delegated pool. In this situation, the client attempts to negotiate IPv6CP with the jpppd process.

By default, when IPCP negotiation is attempted for an IPv4-only PPPoE subscriber session on a dynamic interface, the jpppd process issues a Protocol-Reject message if AAA does not provide an IPv4 address. However, negotiation is allowed to proceed when the `on-demand-ip-address` statement is included at the `[edit protocols ppp-service]` or `[edit dynamic-profiles profile-name interfaces pp0 unit $junos-interface-unit ppp-options]` hierarchy level.

IPCP negotiation is enabled by default for an IP destination address defined on a static interface.

In contrast, IPv6CP negotiation is enabled to proceed by default for an IPv6-only session when AAA has not provided an appropriate IPv6 address or prefix. To prevent endless client negotiation of IPv6CP, you can alter the behavior by including the `reject-unauthorized-ipv6cp` statement at the `[edit protocols ppp-service]` hierarchy level. This statement enables the jpppd process to reject the negotiation attempt.

When IPv6CP rejection is enabled, jpppd also issues a Protocol-Reject message when router advertisement is not enabled in the dynamic profile that instantiates the interface but only a Framed-IPv6-Prefix attribute is received.

Shared IPv4 and IPv6 Service Sessions on PPP Access Networks

IN THIS SECTION

- [Accounting for Shared IPv4 and IPv6 Service Sessions | 1328](#)
- [Deactivating Shared IPv4 and IPv6 Service Sessions | 1328](#)

You can configure one dynamic service profile that supports IPv4, IPv6, or both IPv4 and IPv6. It allows subscribers to share the same service session using IPv4 and IPv6 address families. If you define IPv4 and IPv6 in the dynamic service profile, one address family or both address families can be activated for the service. When the service is activated, matched packets are tagged with the same traffic class and treated the same way for both IPv4 and IPv6 traffic.

Accounting for Shared IPv4 and IPv6 Service Sessions

When service sessions are shared for both IPv4 and IPv6 subscribers, only one Accounting-Start message is sent for each service session regardless of the number of address families that are active. Statistics for each address family of a service session are cumulative across service activations and deactivations of the service.

Deactivating Shared IPv4 and IPv6 Service Sessions

If both IPv4 and IPv6 service sessions are active, and a deactivation message is received for one of the address families (IPv4 or IPv6), all active services for that address family are deactivated. If one address family remains active on the service, the service session remains in the ACTIVE state. If the address family that is deactivated is the only family currently running on the service session, the service returns to the INIT state.

AAA Service Framework in a Dual Stack over a PPPoE Access Network

IN THIS SECTION

- [Collection of Accounting Statistics in a PPPoE Access Network | 1330](#)

- [Change of Authorization \(CoA\) | 1331](#)

You can use the AAA Service Framework for all authentication, authorization, accounting, address assignment, and dynamic request services that the BNG uses for network access. The framework supports authentication and authorization through external RADIUS servers. It also supports accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS servers.

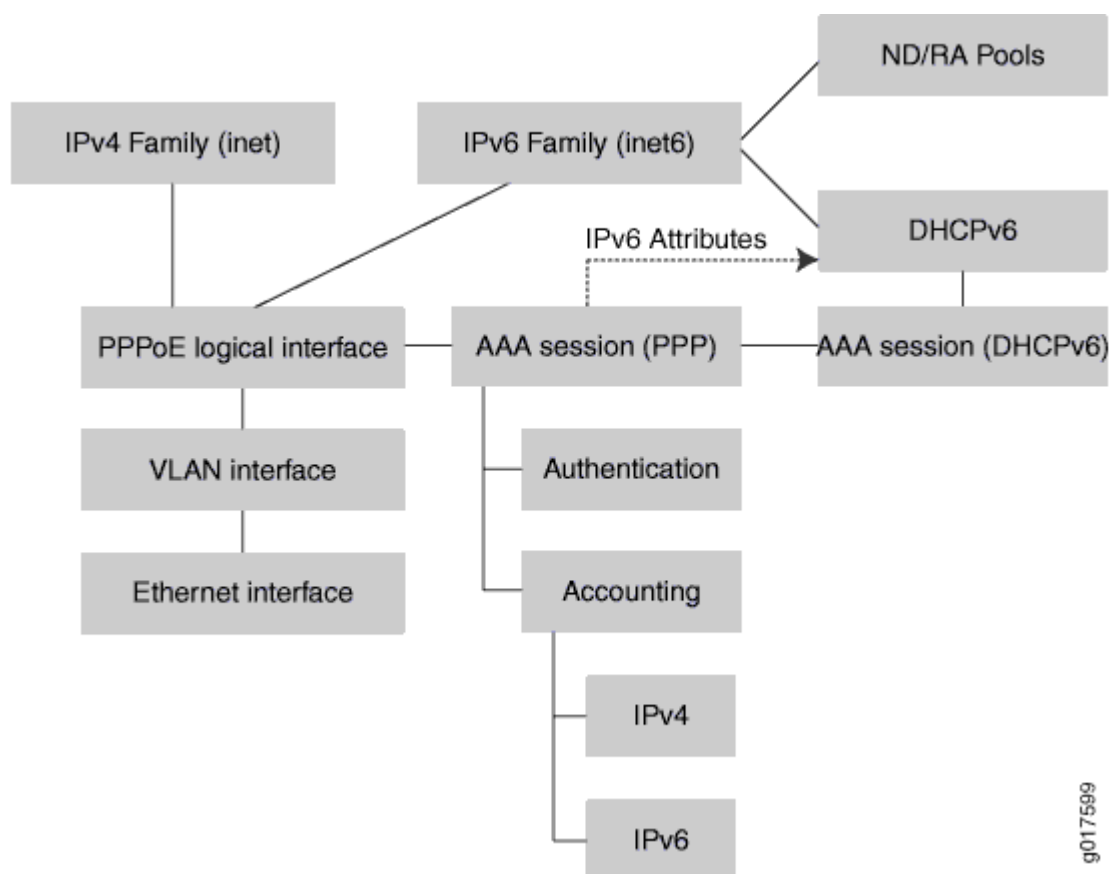
The BNG interacts with external servers to determine how individual subscribers access the broadband network. The BNG can also obtain information from external servers for the following:

- How subscribers are authenticated.
- How accounting statistics are collected and used.
- How dynamic requests, such as CoA, are handled.

As shown in [Figure 23 on page 1330](#), implementing a dual stack over a PPPoE access network that uses AAA can have the following characteristics:

- DHCPv6—If used, it runs over the IPv6 family session, and it inherits attributes from the underlying PPPoE session.
- NDRA—If used, it runs over the IPv6 family session.
- IPv4 and IPv6 accounting—One accounting session handles both IPv4 and IPv6 accounting information.

Figure 23: AAA Service Framework in a Dual Stack over a PPPoE Access Network



Collection of Accounting Statistics in a PPPoE Access Network

AAA provides support for both IPv4 and IPv6 statistics in one accounting session. On MX Series 5G Universal Routing Platforms, AAA also provides support for separate IPv4 and IPv6 accounting statistics.

The following RADIUS attributes are included by default (when available) in Acct-Start, Interim, and Acct-Stop messages:

- Framed-IPv6-Prefix
- Framed-IPv6-Pool
- Delegated-Ipv6-Prefix
- Framed-IPv4-Route
- Framed-IPv6-Route

You can configure the BNG to exclude these attributes in Acct-Start and Acct-Stop messages.

Change of Authorization (CoA)

RADIUS servers can initiate dynamic requests to the BNG. Dynamic requests include CoA requests, which specify vendor-specific attribute (VSA) modifications and service changes.

In your access profile configuration, you specify the IP addresses of RADIUS authentication servers that can initiate dynamic requests to the router. The list of authentication servers also provides RADIUS-based dynamic service activation and deactivation during subscriber login.

RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers

Acct-Start messages sent to the RADIUS server contain all the learned and allocated addresses. Subsequent negotiation or allocation of addresses results in optionally sending immediate Acct-Interim-Update messages that contain all the negotiated and allocated addresses. For the dual-stack PPPoE subscriber, the following types of addresses are provided:

- IP address–negotiated during the IPCP (NCP) phase of PPP
- Interface identifier–negotiated during the IPv6CP (NCP) phase of PPP
- NDRA prefix–sent during router advertisement after IPv6CP
- DHCPv6 IA_NA address–negotiated by the DHCPv6 Solicit, Advertise, Request, Reply (SARR) phase after IPv6CP
- DHCPv6 IA_PD prefix–negotiated by the DHCPv6 SARR phase after IPv6CP

The BNG identifies addresses by the following methods:

- Addresses or prefixes returned from an external authority, such as RADIUS
- Addresses allocated locally using the pool names specified by external authority
- Addresses allocated from a local pool not specified for PPP authorization
- Addresses allocated by an external server outside of the BNG or RADIUS, such as a DHCPv6 external server (DHCPv6 relay or relay proxy)

IPCP and IPv6CP negotiation occur at the PPP NCP phase and can occur in any order. However, DHCPv6 PD or DHCPv6 IA_NA allocation and negotiation occur only after IPv6CP.

The following table lists the RADIUS attributes and their mapping:

Number	RADIUS Attribute	Address Type
1	Framed-IP-Address	IP Address

(Continued)

Number	RADIUS Attribute	Address Type
2	Framed-Pool	IP Address Pool
3	Framed-IPv6-Prefix	NDRA_Prefix (prefix < 128) IA_NA (prefix = 128)
4	Framed-IPv6-Pool	NDRA Prefix pool IA_NA pool
5	Framed-Interface-Id	IPv6 Interface Identifier
6	Delegated-IPv6-Prefix	IA_PD Prefix
7	Jnpr-Delegated-IPv6-Pool (VSA 26-161)	IA_PD Pool
8	Jnpr-IPv6-Ndra-Pool-Name (VSA 26-157) NOTE. Not supported: Use Framed-IPv6-Pool to specify the NDRA pool. Alternatively, configure it locally by using the neighbor-discovery-router-advertisement pool statement.	NDRA Pool

Accounting Messages for PPPoE Using NDRA Prefixes

In the following tables, you can compare PPPoE dual-stack address allocation using Stateless Address Autoconfiguration (SLAAC) NDRA.

The following table lists SLAAC (NDRA) prefixes from RADIUS:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	Framed-IPv6-Prefix (used for NDRA Prefix) Framed-Interface-Id Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP	Because the required attributes are learned prior to the Acct-Start messages, these attributes are sent in Acct-Start messages and no immediate Acct-Interim-Update message is sent.
2	Framed-IPv6-Prefix (used for NDRA Prefix) Framed-Interface-Id Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Because the required attributes are learned prior to the Acct-Start messages, these attributes are sent in Acct-Start messages and no immediate Acct-Interim-Update message is sent. No immediate Acct-Interim-Update message is sent after DHCPv6.
3	Framed-IPv6-Prefix (used for NDRA Prefix) Framed-Interface-Id not sent Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message contains only iFramed-IPv6-Prefix and Delegated-IPv6-Prefix. No immediate Acct-Interim-Update message is sent. Next periodic Acct-Interim-Update message (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	<p>Framed-IPv6-Prefix (used for NDRA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix and Delegated-IPv6-Prefix.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Next periodic Acct-Interim-Update message (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.</p>

The following table lists prefixes from RADIUS selected pools:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id</p> <p>Jnpr-Delegated-IPv6-pool (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id.</p> <p>Framed-IPv6-Prefix is based on the configuration present in the dynamic profile IPv6 prefix that was allocated and sent in Acct-Start message.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS.</p> <p>No Acct-Interim-Update message is sent.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id.</p> <p>Framed-IPv6-Prefix and Framed-IPv6-Pool are based on the configuration present in dynamic profile IPv6 prefix and is allocated prior and sent in Acct-Start message.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>No immediate Acct-Interim-Update message is sent.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. (This value is learned during IPv6NCP negotiation with the peer.)</p>
4	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. (This value is learned during IPv6NCP negotiation with the peer.)</p>

The following table lists prefixes from a local pool or an external server:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool.</p> <p>No immediate Acct-Interim-Update message is sent.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Framed-IPv6-Pool.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	<p>Framed-IPv6-Pool (used for NDRA Prefix)</p> <p>NOTE: If RADIUS does not return Framed-IPv6-Pool, you can configure this locally using the neighbor-discovery-router-advertisement pool statement, which is used for allocating an NDRA prefix from the local pool.</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p> <p>[IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Upon DHCPv6, immediate Acct-Interim-Update is sent which contains Framed-IPv6-Pool, Framed-IPv6-Prefix, Framed-Interface-Id, and DHCPv6 IA_PD Prefix.</p> <p>Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer.</p> <p>DHCPv6 IA_PD is learned by DHCPv6 either by DHCPv6 external server or reservation from a local pool during DHCPv6 SARR phase.</p> <p>Any periodic Acct-Interim-Update before DHCPv6 completion contains Framed-Interface-Id in addition to the attributes of the Acct-Start message. (This can occur if DHCPv6 occurs after periodic interval.)</p>

Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes

In the following tables, you can compare PPPoE dual-stack address allocation using DHCPv6 IA_NA prefixes.

The following table lists DHCPv6 IA_NA prefixes from RADIUS:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	<p>Framed-IPv6-Prefix (used for IA_NA prefix)</p> <p>Framed-Interface-Id</p> <p>Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Because the required attributes are learned prior to Acct-Start message, these attributes are sent in Acct-Start message and no immediate Acct-Interim-Update message is sent.</p>
2	<p>Framed-IPv6-Prefix (used for IA_NA Prefix)</p> <p>Framed-Interface-Id</p> <p>Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Because the required attributes are learned prior to Acct-Start message, these attributes are sent in Acct-Start message and no immediate Acct-Interim-Update message is sent.</p> <p>There is no immediate Acct-Interim-Update message sent after DHCPv6.</p>
3	<p>Framed-IPv6-Prefix (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Acct-Start message message contains Framed-IPv6-Prefix and Delegated-IPv6-Prefix.</p> <p>No immediate Acct-Interim-Update message is sent.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	<p>Framed-IPv6-Prefix (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains iFramed-IPv6-Prefix and Delegated-IPv6-Prefix.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.</p>

The following table lists prefixes from RADIUS selected pools:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	<p>Framed-IPv6-Pool (used for IA_NA Prefix)</p> <p>Framed-Interface-Id</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id .</p> <p>Framed-IPv6 Prefix is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>Delegated-IPv6-Prefix is pre-allocated .</p> <p>Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS.</p> <p>No Acct-Interim-Update message is sent.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	<p>Framed-IPv6-Pool (used for IA_NAPrefix)</p> <p>Framed-Interface-Id</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id.</p> <p>Framed-IPv6 is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	<p>Framed-IPv6-Pool (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix.</p> <p>Framed-IPv6-Prefix is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>No immediate Acct-Interim-Update message is sent.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. (This value is learned during IPv6NCP negotiation with the peer.)</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	<p>Framed-IPv6-Pool (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix.</p> <p>Framed-IPv6-Prefix is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>Delegated-IPv6-Prefix is pre-allocated.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. (This value is learned during IPv6NCP negotiation with the peer.)</p>

The following table lists prefixes from a local pool or an external server:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	<p>Framed-IPv6-Pool (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p>	IPv6NCP	<p>Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool.</p> <p>Framed-IPv6-Prefix is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>No immediate Acct-Interim-Update message is sent.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Framed-IPv6-Pool. (This value is learned during IPv6NCP negotiation with the peer.)</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	<p>Framed-IPv6-Pool (used for IA_NA Prefix)</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p> <p>[IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p>	IPv6NCP + DHCPv6	<p>Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool.</p> <p>Framed-IPv6-Prefix is pre-allocated.</p> <p>Framed-IPv6-Pool is learned from RADIUS.</p> <p>No immediate Acct-Interim-Update message is sent upon DHCPv6.</p> <p>Upon DHCPv6, an immediate Acct-Interim-Update is sent that contains Framed-IPv6-Pool, Framed-IPv6-Prefix, Framed-Interface-Id, and DHCPv6 IA_PD Prefix.</p> <p>Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer.</p> <p>DHCPv6 IA_PD is learned by DHCPv6 either by DHCPv6 External Server or reservation from a local pool during DHCPv6 SARR phase.</p> <p>Any periodic Acct-Interim-Update before DHCPv6 completion contains Framed-Interface-Id in addition to the attributes of Acct-Start message. (This can occur if DHCPv6 occurs after periodic interval.)</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	<p>Framed-IPv6-Pool not sent</p> <p>[IA_NA Prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p> <p>[IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p>	IPv6NCP	<p>Acct-Start message does not contain any of the IPv6-related prefixes and addresses or pool names.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id. (This value is learned during IPv6NCP negotiation with the peer.)</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	<p>Framed-IPv6-Pool not sent</p> <p>[IA_NA Prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p> <p>Framed-Interface-Id not sent</p> <p>Jnpr-Delegated-IPv6-Pool not sent</p> <p>[IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]</p>	IPv6NCP+DHCPv6	<p>Acct-Start message does not contain any of the IPv6-related prefixes and addresses or pool names.</p> <p>No immediate Acct-Interim-Update message is sent upon IPv6NCP.</p> <p>Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id. (This value is learned during IPv6NCP negotiation with the peer.)</p> <p>Upon DHCPv6, immediate Acct-Interim-Update is sent which contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Framed-Interface-Id, and DHCPv6 IA_PD.</p> <p>Framed-IPv6-Prefix is the IA_NA prefix learned by DHCPv6 (either by external server or reservation from a local pool).</p> <p>Framed-IPv6-Pool is sent only if there is a reservation of an IA_NA prefix from local pool by DHCPv6.</p> <p>Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer.</p> <p>DHCPv6 IA_PD prefix is learned by DHCPv6 either by DHCPv6 external server or reservation from a local pool during DHCPv6 SARR phase.</p> <p>Any periodic Acct-Interim-Update before DHCPv6 completion contains Framed-Interface-Id only.</p>

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
			(This can occur if DHCPv6 occurs after periodic interval.)

Suppressing Accounting Information That Comes from AAA

The following standard and vendor-specific IPv6 RADIUS attributes are included by default (when available) in Acct-Start and Acct-Stop messages:

- Framed-IPv6-Prefix
- Framed-IPv6-Pool
- Delegated-Ipv6-Prefix
- Framed-IPv4-Route
- Framed-IPv6-Route

You can configure the software to exclude these attributes from Acct-Start or Acct-Stop messages. To do so, configure the access profile:

1. Access the access profile.

```
[edit]
user@host# edit access profile dual-stack radius attributes
```

2. The following examples show how to use the `exclude` statement to exclude attributes from messages.

```
[edit access profile dual-stack radius attributes]
user@host# set exclude delegated-ipv6-prefix accounting-start
```

```
[edit access profile dual-stack radius attributes]
user@host# set exclude framed-ipv6-pool [accounting-start accounting-stop]
```

```
[edit access profile dual-stack radius attributes]
user@host# set exclude framed-ipv6-prefix accounting-start framed-ipv6-route accounting-start
```

```
[edit access profile dual-stack radius attributes]
user@host# set exclude framed-ipv6-prefix accounting-start framed-ipv6-route accounting-start
```

Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address

You can control the behavior of the router in a situation where IPv6CP negotiation is initiated for subscriber sessions when no authorized addresses are available.

By default, IPv6CP negotiation is enabled to proceed for an IPv6-only session when AAA has not provided an appropriate IPv6 address or prefix. In the absence of the address, the negotiation cannot successfully complete. To prevent endless client negotiation of IPv6CP, include the `reject-unauthorized-ipv6cp` statement at the `[edit protocols ppp-service]` hierarchy level, which enables the `jpppd` process to reject the negotiation attempt.

To configure the router to reject IPv6CP negotiation messages when no IPv6 address is available for a dynamic interface:

- Enable rejection of unauthorized IPv6CP negotiation messages.

```
[edit protocols ppp-service]
user@host# set reject-unauthorized-ipv6cp
```



NOTE: The `reject-unauthorized-ipv6cp` statement does not prevent IPv6CP negotiation for static interfaces, because the `jpppd` process cannot determine whether router advertisement of DHCPv6 is configured to run above the PPP interface.

RELATED DOCUMENTATION

[Migration to IPv6 Using IPv4 and IPv6 Dual Stack | 1245](#)

[Dual Stack for PPPoE Access Networks Using DHCP | 1356](#)

[Dual Stack for PPPoE Access Networks Using NDRA | 1360](#)

[Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | 1425](#)

Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network

IN THIS SECTION

- [Best Practice: Static PPPoE Interfaces with NDRA | 1351](#)
- [Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network | 1352](#)
- [Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | 1353](#)
- [Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 | 1353](#)
- [Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles | 1354](#)
- [Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | 1355](#)

Best Practice: Static PPPoE Interfaces with NDRA

When you use static PPPoE interfaces with NDRA, the prefix configured for router advertisement must match the source address specified under family inet6 in the logical pp0 interface configuration. If these values do not match, the prefix is not advertised correctly.

For example:

```
[edit protocols router-advertisement]
interface pp0.2004 {
```

```
prefix 2001:db8:2040:2004::/64;
}
```

```
[edit interface pp0]
unit 2004 {
  family inet6 {
    address 2001:db8:2040:2004::10.1.1.1/64;
  }
}
```

To view the prefix in the ICMPv6 packet, use the `monitor traffic interface pp0.xxx extensive` command. If the prefix is missing, make sure that there is not a mismatch between the family inet6 address configured for the interface and the prefix configured for the interface in the router advertisement configuration.

Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network

When you use DHCPv6 prefix delegation over a PPPoE access network, you need to enable unnumbered addressing in the family inet6 configuration.

For dynamic PPPoE interfaces, enable unnumbered addressing in the dynamic profile. For example:

```
[edit dynamic-profiles]
PPPoE-dyn-ipv4v6-dhcp {
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        . . .
        family inet6 {
          unnumbered-address 100.0;
        }
      }
    }
  }
}
```

For static PPPoE interfaces, enable unnumbered addressing in the interface configuration. For example:

```
[edit interface pp0]
unit 2004 {
```

```
family inet6 {
  unnumbered-address lo0.0;
```

Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA

When you use NDRA, always set the IPv6 internet address in dynamic profiles to the `$junos-ipv6-address` predefined variable. This variable is replaced with the IPv6 address of the interface used for router advertisements.

```
[edit dynamic-profiles]
dyn-v4v6-ndra {
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        family inet6 {
          address "$junos-ipv6-address ";
        }
      }
    }
  }
}
```

Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6

The IPv6 address configuration for logical interfaces in PPPoE dynamic profiles when you are using DHCPv6 depends on whether or not you are using routing instances.

If you are using routing instances, use the `$junos-loopback-interface` predefined variable for the IPv6 address. For example:

```
[edit dynamic-profiles]
dyn-v4v6-ri {
  routing-instances {
    "$junos-routing-instance" {
      interface "$junos-interface-name";
    }
  }
  interfaces {
    pp0 {
```



```

        unit "$junos-interface-unit" {
            family inet6 {
                unnumbered-address "$junos-loopback-interface";
            }
        }
    }
}

```

If you are not using routing instances, use the unnumbered address for the IPv6 address. The unnumbered address enables the local address to be derived from the specified interface and allows IP processing on the interface without assigning an explicit IP address to the interface. For example:

```

[edit dynamic-profiles]
dyn-v4v6-ndra {
    interfaces {
        pp0 {
            unit "$junos-interface-unit" {
                pppoe-options {
                    underlying-interface "$junos-underlying-interface";
                    server;
                }
                family inet6 {
                    unnumbered-address lo0.0;
                }
            }
        }
    }
}

```

Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles

The IPv4 address configuration for logical interfaces in PPPoE dynamic profiles depends on whether or not you are using routing instances.

If you are using routing instances, use the `$junos-loopback-interface` variable for the IPv6 address.

```

[edit dynamic-profiles]
dyn-v4v6-ri {
    routing-instances {
        "$junos-routing-instance" {

```

```

        interface "$junos-interface-name";
    }
}
interfaces {
    pp0 {
        unit "$junos-interface-unit" {
            family inet {
                unnumbered-address "$junos-loopback-interface";
            }
        }
    }
}
}

```

If you are not using routing instances, use the unnumbered address for the IPv6 address. The unnumbered address enables the local address to be derived from the specified interface and allows IP processing on the interface without assigning an explicit IP address to the interface.

```

[edit dynamic-profiles]
dyn-v4v6-ndra {
    interfaces {
        pp0 {
            unit "$junos-interface-unit" {
                pppoe-options {
                    underlying-interface "$junos-underlying-interface";
                    server;
                }
                family inet {
                    unnumbered-address 100.0;
                }
            }
        }
    }
}
}

```

Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network

In most cases PPPoE is used to authenticate subscribers in a PPPoE access network. However, if you wish to use DHCP to perform the authentication, do not configure subscriber authentication at the [edit system services dhcp-local-server] or the [edit system services dhcp-local-server dhcpv6] hierarchy levels.

Instead configure subscriber authentication at the [edit system services dhcp-local-server dhcpv6 group] hierarchy level. For example:

```
[edit system services dhcp-local-server dhcpv6]
group v6-dhcp-client {
  authentication {
    password $ABC123;
    username-include {
      user-prefix StaticUser;
    }
  }
}
```

RELATED DOCUMENTATION

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Dual Stack for PPPoE Access Networks Using DHCP | 1356](#)

[Dual Stack for PPPoE Access Networks Using NDRA | 1360](#)

Dual Stack for PPPoE Access Networks Using DHCP

IN THIS SECTION

- [Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | 1356](#)
- [Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network | 1357](#)

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE

To layer DHCPv6 above the PPPoE IPv6 family (inet6), create a DHCPv6 local server and associate DHCPv6 with the PPPoE interfaces by adding the PPPoE interfaces to the DHCPv6 local server configuration. Specify static and dynamic PPPoE interfaces as follows:

- **Dynamic**—Use the pp0.0 (PPPoE) logical interface as a wildcard to indicate that a DHCPv6 binding can be made on top of a PPPoE interface.

- Static—Use unit numbers to explicitly specify static interfaces; for example, pp0.2000.

To configure a DHCPv6 local server:

1. Access the DHCPv6 local server configuration.

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

2. Create a group for dynamic PPPoE interfaces and assign a name.

The group feature groups a set of interfaces and then applies a common DHCP configuration to the named interface group.

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

3. For dynamic PPPoE logical interfaces, add an interface.

```
[edit system services dhcp-local-server dhcpv6 group group-pppoe]
user@host# set interface pp0.0
```

4. For static PPPoE, add a range of interfaces.

```
[edit system services dhcp-local-server dhcpv6 group group-pppoe]
user@host# set interface pp0.2000 upto pp0.2999
```

SEE ALSO

[IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1325](#)

[Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | 1355](#)

Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network

Configure a dynamic profile for IPv4 and IPv6 subscribers that access the network. The dynamic profile defines the attributes of the dynamic PPPoE logical subscriber interface.

To configure a PPPoE dynamic profile for both IPv4 and IPv6 subscribers:

1. Create and name the dynamic profile.

```
[edit]
user@host# edit dynamic profiles PPPoE-dyn-ipv4v6
```

2. If you are using routing instances, add a routing instance to the profile, and add an interface to the routing instance.

- Specify the `$junos-routing-instance` variable for the routing instance. The routing instance variable is dynamically replaced with the routing instance the accessing subscriber uses when connecting to the BNG.
- Specify the `$junos-interface-name` variable for the interface. The interface variable is dynamically replaced with the interface that the accessing subscriber uses when connecting to the BNG.

```
[edit dynamic profiles PPPoE-dyn-ipv4v6]
user@host# set routing-instances $junos-routing-instance interface $junos-interface-name
```

3. Add a PPPoE logical interface (pp0) to the profile, and specify `$junos-interface-unit` as the predefined variable to represent the logical unit number for the interface. The variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

```
[edit dynamic profiles PPPoE-dyn-ipv4v6]
user@host# edit interfaces pp0 unit $junos-interface-unit
```

4. Configure the IPv4 family for the pp0 interface as follows:

- If you are not using routing instances, assign an unnumbered address. The unnumbered address enables the local address to be derived from the specified interface and allows IP processing on the interface without an explicit IP address assigned to the interface.

For example:

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet unnumbered-address lo0.0
```

- If you are using routing instances, assign the predefined variable `$junos-loopback-interface`.

For example:

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet unnumbered-address $junos-loopback-interface
```

5. Configure the IPv6 family for the pp0 interface as follows:

- If you are not using routing instances, assign an unnumbered address that specifies the loopback interface. The unnumbered address enables the local address to be derived from the loopback interface and allows IP processing on the interface without an explicit IP address assigned to the interface.

For example:

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 unnumbered-address lo0.0
```

- If you are using routing instances, assign the predefined variable \$junos-loopback-interface.

For example:

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 unnumbered-address $junos-loopback-interface
```

6. Specify \$junos-underlying-interface as the predefined variable to represent the name of the underlying Ethernet interface on which the router creates the dynamic PPPoE logical interface. The variable is dynamically replaced with the actual name of the underlying interface.

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

7. Define the router to act as a PPPoE server when a PPPoE logical interface is dynamically created.

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

8. (Optional) Configure the PPP authentication protocol for the pp0 interface. Specify either chap or pap (or both).

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
user@host# set ppp-options pap
```

9. (Optional) Enable keepalives and set an interval for keepalives. We recommend an interval of 30 seconds. For example:

```
[edit dynamic-profiles PPPoE-dyn-ipv4v6 interfaces pp0 unit "$junos-interface-unit"]
user@host# set keepalives interval 30
```

SEE ALSO

[Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 | 1353](#)

RELATED DOCUMENTATION

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1351](#)

Dual Stack for PPPoE Access Networks Using NDRA

IN THIS SECTION

- [Configuring a Static PPPoE Logical Interface for NDRA | 1361](#)
- [Configuring an Address-Assignment Pool Used for Router Advertisements | 1362](#)
- [Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | 1363](#)
- [Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | 1364](#)
- [Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | 1364](#)

- [Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE](#) | 1390

Configuring a Static PPPoE Logical Interface for NDRA

To configure a static PPPoE logical interface for static Neighbor Discovery Router Advertisement (NDRA) configurations:

1. Specify the name and logical unit number of the interface.

```
[edit]
user@host# edit interfaces pp0 unit 1000
```

2. Configure a description for the interface.

```
[edit interfaces pp0 unit 1000]
user@host# set description "static IPv4v6 dual stack, NDRA"
```

3. Specify the family inet6 source address.

```
[edit interfaces pp0 unit 1000]
user@host# set family inet6 address 2001:db8:2040:2004::10.1.1.1/64
```

4. Configure an unnumbered address for family inet.

```
[edit interfaces pp0 unit 1000]
user@host# set family inet unnumbered-address lo0.0
```

5. Specify the underlying Ethernet interface.

```
[edit interfaces pp0 unit 1000]
user@host# set pppoe-options underlying-interface ge-1/0/0.1000
```

6. Define the router to act as a PPPoE server when the PPPoE logical interface is created.

```
[edit interfaces pp0 unit 1000]
user@host# set pppoe-options server
```


7. Access the router advertisement configuration, and specify the prefixes that the BNG sends in router advertisements for the static interface. Make sure that the prefixes match the source address configured for the static PPPoE logical interface configured in Step 3.

```
[edit]
user@host# edit protocols router-advertisement
user@host# set interface pp0.1000 prefix 2001:db8:2040:2004::/64
```

SEE ALSO

[Best Practice: Static PPPoE Interfaces with NDRA | 1351](#)

Configuring an Address-Assignment Pool Used for Router Advertisements

If you are using local address-assignment pools to be used for router advertisement, create a pool and add IPv6 prefixes to the pool.

You must configure separate pools for DHCPv6 prefix delegation, DHCPv6 IA_NA, and router advertisement.

To configure an NDRA address-assignment pool.

1. Create a pool for IPv6 prefixes used by NDRA.

```
[edit]
user@host# edit access address-assignment pool ndra-2010 family inet6
```

2. Add IPv6 network prefixes to the pool.

```
[edit access address-assignment pool ndra-2010 family inet6]
user@host# set prefix 2001:db8::/64
```

3. Configure the name of the IPv6 address range and define the range. For NDRA pools, specify the range by setting a prefix length of 64.

```
[edit access address-assignment pool ndra-2010 family inet6]
user@host# set range ndra-range prefix-length 64
```

4. Specify that the address-assignment pool is used for NDRA.

```
[edit access address-assignment]
user@host# set neighbor-discovery-router-advertisement ndra-2010
```

SEE ALSO

[Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250](#)

[Methods for Obtaining IPv6 Prefixes for NDRA | 1255](#)

Configuring Duplicate IPv6 Prefix Protection for Router Advertisement

If you are using AAA to supply IPv6 prefixes for router advertisement, you can enable duplicate prefix protection to prevent prefixes from being used more than once. If enabled, the following attributes received from external servers are checked:

- *Framed-IPv6-Prefix*
- *Framed-IPv6-Pool*

The router then takes one of the following actions:

- If a prefix matches a prefix in an address pool, the prefix is taken from the pool if it is available.
- If the prefix is already in use, it is rejected as unavailable.
- If the prefix length requested from the external server does not match the pool's prefix length exactly, the authentication request is denied. If configured, the Acct-Stop message will include a termination cause.

To configure duplicate prefix protection:

1. Enter the access configuration.

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

2. Enable duplicate prefix protection.

```
[edit access]
user@host# address-protection
```

SEE ALSO

[Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250](#)

[Duplicate Prefix Protection for NDRA | 1256](#)

Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces

When you are using Router Advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces, both interfaces should use the same link-local address.

The link local address should be assigned using a unique 64-Bit IPv6 interface identifier (EUI-64), which is obtained based on the MAC address of the underlying interface.

To cause the system to implement the link-local address based on the MAC address of the underlying interface and to comply with the 64-bit Extended Unique Identifier (EUI-64):

1. Access the hierarchy that configures all static demux interfaces on the router.

```
[edit]
edit system demux-options
```

2. Configure the system to use the MAC address of the underlying static interface as the basis for the link-local address of the demux interface.

```
[edit system demux-options]
set use-underlying-interface-mac
```

Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE

IN THIS SECTION

- [Requirements | 1365](#)
- [Overview | 1365](#)
- [Configuration | 1366](#)
- [Verification | 1384](#)

This example shows a dual stack configuration for a residential subscriber with a single PC. It uses ND/RA to provide a prefix used to obtain a global IPv6 address for the PC.

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router
- Junos OS Release 11.4 or later

Overview

IN THIS SECTION


- [Topology](#) | 1365

This design uses ND/RA in your subscriber access network as follows:

- The access network is PPPoE.
- ND/RA is used to assign a global IPv6 address on the WAN link. The prefixes used in router advertisements come from a local pool that is specified using AAA RADIUS.

Topology

Figure 24: PPPoE Subscriber Access Network with NDRA



g017769

Table 59 on page 1366 describes the configuration components used in this example.

Table 59: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation

Configuration Component	Component Name	Purpose
Dynamic Profiles	DS-dyn-ipv4v6-ndra	Profile that creates a PPPoE logical interface when the subscriber logs in.
Interfaces	ge-3/3/0	Underlying Ethernet interface.
	lo0	Loopback interface for use in the access network. The loopback interface is automatically used for unnumbered interfaces.
Address-Assignment Pools	default-ipv4-pool-2	Pool that provides IPv4 addresses for the subscriber LAN.
	ndra-2010	Pool that provides IPv6 prefixes used in router advertisements. These prefixes are used to create a global IPv6 address that is assigned to the CPE WAN link.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1367](#)
- [Configuring a Dynamic Profile for the PPPoE Logical Interface | 1370](#)
- [Configuring a Loopback Interface | 1373](#)
- [Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces | 1375](#)
- [Specifying the BNG IP Address | 1377](#)
- [Configuring RADIUS Server Access | 1378](#)
- [Configuring RADIUS Server Access Profile | 1380](#)
- [Specifying the RADIUS Server Access Profile to Use | 1381](#)
- [Configuring Local Address-Assignment Pools | 1382](#)

To configure this example, perform these tasks:

CLI Quick Configuration

The following is the complete configuration for this example:

```
dynamic-profiles {
  DS-dyn-ipv4v6-ra {
    interfaces {
      pp0 {
        unit "$junos-interface-unit" {
          ppp-options {
            chap;
            pap;
          }
          pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
          }
          keepalives interval 30;
          family inet {
            unnumbered-address lo0.0;
          }
          family inet6 {
            address $junos-ipv6-address;
          }
        }
      }
    }
  }
  protocols {
    router-advertisement {
      interface "$junos-interface-name" {
        prefix $junos-ipv6-ndra-prefix;
      }
    }
  }
}
system {
  services {
    dhcp-local-server {
      dhcpv6 {
```

```

        group DHCPv6-over-pppoe {
            interface pp0.0;
        }
    }
}

access-profile Access-Profile;
interfaces {
    ge-3/3/0 {
        unit 1004 {
            description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
            encapsulation ppp-over-ether;
            vlan-id 1004;
            pppoe-underlying-options {
                duplicate-protection;
                dynamic-profile DS-dyn-ipv4v6-ra;
            }
        }
    }
    lo0 {
        description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
        unit 0 {
            family inet {
                address 192.0.2.77/32 {
                    primary;
                }
            }
            family inet6 {
                address 2001:db8:2030:0:0:0::1/64 {
                    primary;
                }
            }
        }
    }
}

routing-options {
    router-id 203.0.113.0;
}

access {
    radius-server {
        203.0.113.99 {
            secret "$ABC123$ABC123ABC123"; ## SECRET-DATA
        }
    }
}

```

```

        timeout 45;
        retry 4;
        source-address 203.0.113.1;
    }
}
profile Access-Profile {
    authentication-order radius;
    radius {
        authentication-server 203.0.113.99;
        accounting-server 203.0.113.99;
    }
    accounting {
        order [ radius none ];
        update-interval 120;
        statistics volume-time;
    }
}
address-assignment {
    neighbor-discovery-router-advertisement ndra-2010;
    pool default-ipv4-pool-2 {
        family inet {
            network 203.0.113.10/16;
            range r5 {
                low 203.0.113.11;
                high 203.0.113.150;
            }
        }
    }
    pool ndra-2010 {
        family inet6 {
            prefix 2001:db8:2010:0:0:0::/48;
            range L prefix-length 64;
        }
    }
}
address-protection;
}

```


Configuring a Dynamic Profile for the PPPoE Logical Interface

CLI Quick Configuration

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

```
edit dynamic-profiles DS-dyn-ipv4v6-ra
edit interfaces pp0 unit $junos-interface-unit
set family inet unnumbered-address lo0.0
set family inet6 address $junos-ipv6-address
set pppoe-options underlying-interface "$junos-underlying-interface"
set pppoe-options server
set ppp-options pap
set ppp-options chap
set keepalives interval 30
up 3
edit protocols router-advertisement
edit interface $junos-interface-name
set prefix $junos-ipv6-ndra-prefix
```

Step-by-Step Procedure

Create a dynamic profile for the PPPoE logical interface. This dynamic profile supports both IPv4 and IPv6 sessions on the same logical interface.

To configure the dynamic profile:

1. Create and name the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles DS-dyn-ipv4v6-ra
```

2. Configure a PPPoE logical interface (pp0) that is used to create logical PPPoE interfaces for the IPv4 and IPv6 subscribers.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# edit interfaces pp0
```

3. Specify `$junos-interface-unit` as the predefined variable to represent the logical unit number for the `pp0` interface. The variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0]
user@host# edit unit $junos-interface-unit
```

4. Specify `$junos-underlying-interface` as the predefined variable to represent the name of the underlying Ethernet interface on which the router creates the dynamic PPPoE logical interface. The variable is dynamically replaced with the actual name of the underlying interface supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

5. Configure the router to act as a PPPoE server when a PPPoE logical interface is dynamically created.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

6. Configure the IPv4 family for the `pp0` interface. Specify the unnumbered address to dynamically create loopback interfaces.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet unnumbered-address lo0.0
```

7. Configure the IPv6 family for the `pp0` interface. Because the example uses router advertisement, assign the predefined variable `$junos-ipv6-address`.

```
[edit dynamic-profilesDS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 unnumbered-address $junos-ipv6-address
```

8. Configure one or more PPP authentication protocols for the pp0 interface.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
user@host# set ppp-options pap
```

9. Enable keepalives and set an interval for keepalives. We recommend an interval of 30 seconds.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set keepalives interval 30
```

10. Access the router advertisement configuration.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# edit protocols router-advertisement
```

11. Specify the interface on which the ND/RA configuration is applied.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra protocols router-advertisement]
user@host# edit interface $junos-interface-name
```

12. Specify a prefix value contained in router advertisement messages sent to the CPE on interfaces created with this dynamic profile. If you specify the \$junos-ipv6-ndra-prefix predefined variable, the actual value is obtained from a local pool or through AAA.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra protocols router-advertisement interface "$junos-
interface-name"]
user@host# set prefix $junos-ipv6-ndra-prefix
```

Results

From configuration mode, confirm your configuration by entering the show command.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# show
interfaces {
  pp0 {
```

```

    unit "$junos-interface-unit" {
        ppp-options {
            chap;
            pap;
        }
        pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
        }
        keepalives interval 30;
        family inet {
            unnumbered-address lo0.0;
        }
        family inet6 {
            address $junos-ipv6-address;
        }
    }
}

protocols {
    router-advertisement {
        interface "$junos-interface-name" {
            prefix $junos-ipv6-ndra-prefix;
        }
    }
}

```

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

Configuring a Loopback Interface

CLI Quick Configuration

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

```

edit interfaces lo0 unit 0
set family inet address 192.0.2.77/32 primary
set family inet6 address 2001:db8:2030:0:0::1/64 primary

```

Step-by-Step Procedure

To configure a loopback interface:

1. Create the loopback interface and specify a unit number.

```
[edit]
user@host# edit interfaces lo0 unit 0
```

2. Configure the interface for IPv4.

```
[edit interfaces lo0 unit 0]
user@host# set family inet address 192.0.2.77/32 primary
```

3. Configure the interface for IPv6.

```
[edit interfaces lo0 unit 0]
user@host# set family inet6 address 2001:db8:2030:0:0::1/64 primary
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces lo0]
user@host# show
unit 0 {
  family inet {
    address 192.0.2.77/32 {
      primary;
    }
  }
  family inet6 {
    address 2001:db8:2030:0:0::1/64 {
      primary;
    }
  }
}
```

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

Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces

CLI Quick Configuration

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

```
edit interfaces ge-3/3/0 unit 1004
set description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd"
set encapsulation ppp-over-ether
set vlan-id 1004
set pppoe-underlying-options duplicate-protection
set pppoe-underlying-options dynamic-profile DS-dyn-ipv4v6-ra
```

Step-by-Step Procedure

To configure the underlying Ethernet interface:

1. Specify the name and logical unit number of the static underlying Ethernet interface to which you want to attach the IPv4 and IPv6 dynamic profile.

```
[edit]
user@host# edit interfaces ge-3/3/0 unit 1004
```

2. Configure a description for the interface.

```
[edit interfaces ge-3/3/0 unit 1004]
user@host# set description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd"
```

3. Configure PPPoE encapsulation on the underlying interface.

```
[edit interfaces ge-3/3/0 unit 1004]
user@host# set encapsulation ppp-over-ether
```

4. Configure the VLAN Id.

```
[edit interfaces ge-3/3/0 unit 1004]
user@host# set vlan-id 1004
```

5. Attach the dynamic profile to the underlying interface.

```
[edit interfaces ge-3/3/0 unit 1004]
user@host# set pppoe-underlying-options dynamic-profile DS-dyn-ipv4v6-ra
```

6. (Optional) Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on the same VLAN interface.

```
[edit interfaces ge-3/3/0 unit 1004]
user@host# set pppoe-underlying-options duplicate-protection
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces]
user@host# show
ge-3/3/0 {
  unit 1004 {
    description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
    encapsulation ppp-over-ether;
    vlan-id 1004;
    pppoe-underlying-options {
      duplicate-protection;
      dynamic-profile DS-dyn-ipv4v6-ra;
    }
  }
}
```

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

Specifying the BNG IP Address

CLI Quick Configuration

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

```
edit routing-options
set router-id 203.0.113.0
```



BEST PRACTICE: We strongly recommend that you configure the BNG IP address to avoid unpredictable behavior if the interface address on a loopback interface changes.

Step-by-Step Procedure

To configure the IP address of the BNG:

1. Access the routing-options configuration.

```
[edit]
user@host# edit routing-options
```

2. Specify the IP address or the BNG.

```
[edit routing-options]
user@host# set router-id 203.0.113.0
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit routing-options]
user@host# show
router-id 203.0.113.0;
```

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

Configuring RADIUS Server Access

CLI Quick Configuration

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

```
edit access radius-server 203.0.113.99
set secret "$ABC123$ABC123ABC123"
set timeout 45
set retry 4
set source-address 203.0.113.1
```

Step-by-Step Procedure

To configure RADIUS servers:

1. Create a RADIUS server configuration, and specify the address of the server.

```
[edit]
user@host# edit access radius-server 203.0.113.99
```

2. Configure the required secret (password) for the server. Secrets enclosed in quotation marks can contain spaces.

```
[edit access radius-server 203.0.113.99]
user@host# set secret "$ABC123$ABC123ABC123"
```

3. Configure the source address that the BNG uses when it sends RADIUS requests to the RADIUS server.

```
[edit access radius-server 203.0.113.99]
user@host# set source address 203.0.113.1
```

4. (Optional) Configure the number of times that the router attempts to contact a RADIUS accounting server. You can configure the router to retry from 1 through 16 times. The default setting is 3 retry attempts.

```
[edit access radius-server 203.0.113.99]
user@host# set retry 4
```

5. (Optional) Configure the length of time that the local router or switch waits to receive a response from a RADIUS server. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

```
[edit access radius-server 203.0.113.99]
user@host# set timeout 45
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
radius-server {
  203.0.113.99 {
    secret "$ABC123$ABC123ABC123"; ## SECRET-DATA
    timeout 45;
    retry 4;
    source-address 203.0.113.1;
  }
}
```

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

Configuring RADIUS Server Access Profile

CLI Quick Configuration

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

```
edit access profile Access-Profile
set authentication-order radius
set radius authentication-server 203.0.113.99
set radius accounting-server 203.0.113.99
set accounting order radius
set accounting order none
set accounting update-interval 120
set accounting statistics volume-time
```

Step-by-Step Procedure

To configure a RADIUS server access profile:

1. Create a RADIUS server access profile.

```
[edit]
user@host# edit access profile Access-Profile
```

2. Specify the order in which authentication methods are used.

```
[edit access profile Access-Profile]
user@host# set authentication-order radius
```

3. Specify the address of the RADIUS server used for authentication and the server used for accounting.

```
[edit access profile Access-Profile]
user@host# set radius authentication-server 203.0.113.99
user@host# set radius accounting-server 203.0.113.99
```

4. Configure RADIUS accounting values for the access profile.

```
[edit access profile Access-Profile]
user@host# set accounting order [ radius none ]
user@host# set accounting update-interval 120
user@host# set accounting statistics volume-time
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
profile Access-Profile {
  authentication-order radius;
  radius {
    authentication-server 203.0.113.99;
    accounting-server 203.0.113.99;
  }
  accounting {
    order [ radius none ];
    update-interval 120;
    statistics volume-time;
  }
}
```

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

Specifying the RADIUS Server Access Profile to Use

CLI Quick Configuration

To quickly configure this example, copy the following command and paste it into the CLI at the `[edit]` hierarchy level.

```
set access-profile Access-Profile
```

Step-by-Step Procedure

To specify the RADIUS server access profile to use for authentication:

1. Specify the access profile.

```
[edit]
user@host# set access-profile Access-Profile
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit]
user@host# show
...
access-profile Access-Profile;
...
```

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

Configuring Local Address-Assignment Pools

CLI Quick Configuration

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

```
edit access
set address-assignment pool default-ipv4-pool-2 family inet network 203.0.113.10/16
set address-assignment pool default-ipv4-pool-2 family inet range r5 low 203.0.113.11
set address-assignment pool default-ipv4-pool-2 family inet range r5 high 203.0.113.150
set address-assignment pool ndra-2010 family inet6 prefix 2001:db8:2010:0:0:0::/48
set address-assignment pool ndra-2010 family inet6 range L prefix-length 64
set address-assignment neighbor-discovery-router-advertisement ndra-2010
set address-protection
```

Step-by-Step Procedure

Configure three address-assignment pools for DHCPv4, DHCPv6 prefix delegation, and ND/RA.

To configure the address-assignment pools:

1. Configure the address-assignment pool for DHCPv4.

```
[edit]
user@host# edit access address-assignment pool default-ipv4-pool-2
user@host# edit family inet
user@host# set network 203.0.113.10/16
user@host# set range r5 low 203.0.113.11
user@host# set range r5 high 203.0.113.150
```

2. Configure the address-assignment pool for ND/RA.

```
[edit]
user@host# edit access address-assignment pool ndra-2010
user@host# edit family inet6
user@host# set prefix 2001:db8:2010:0:0:0::/48
user@host# set range L prefix-length 64
```

3. Specify that the address-assignment pool is used for NDRA.

```
[edit]
user@host# edit access address-assignment
user@host# set neighbor-discovery-router-advertisement ndra-2010
```

4. (Optional) Enable duplicate prefix protection.

```
[edit access]
user@host# set address-protection
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
address-assignment {
  neighbor-discovery-router-advertisement ndra-2010;
  pool default-ipv4-pool-2 {
    family inet {
      network 203.0.113.10/16;
      range r5 {
        low 203.0.113.11;
        high 203.0.113.150;
      }
    }
  }
  pool ndra-2010 {
    family inet6 {
      prefix 2001:db8:2010:0:0:0::/48;
      range L prefix-length 64;
    }
  }
}
address-protection;
```

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

Verification

IN THIS SECTION

- [Verifying Active Subscriber Sessions | 1385](#)
- [Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | 1385](#)
- [Verifying Dynamic Subscriber Sessions | 1386](#)
- [Verifying the ND/RA Prefix Pool and Prefix Length | 1387](#)
- [Verifying the Status of the PPPoE Logical Interface | 1388](#)
- [Verifying Router Advertisements | 1389](#)

Confirm that the configuration is working properly.

Verifying Active Subscriber Sessions

Purpose

Verify active subscriber sessions.

Action

From operational mode, enter the `show subscribers summary` command.

```
user@host>show subscribers summary
Subscribers by State
  Active: 2
  Total: 2

Subscribers by Client Type
  DHCP: 1
  PPPoE: 1
  Total: 2
```

Meaning

The fields under `Subscribers by State` show the number of active subscribers.

The fields under `Subscribers by Client Type` show the number of active DHCP and DHCPoE subscriber sessions.

Verifying Both IPv4 and IPv6 Address in Correct Routing Instance

Purpose

Verify that the subscriber has both an IPv4 and IPv6 address and is placed in the correct routing instance.

Action

From operational mode, enter the `show subscribers` command.

```
user@host>show subscribers
Interface      IP Address/VLAN ID  User Name          LS:RI
pp0.1073741864 192.0.2.5           dual-stack-v4v6-pd default:default
*              2001:db8:2010:0:0:8::/64
pp0.1073741864 2001:db8:2040:2000:2000:5::/64          default:default
```

Meaning

The Interface field shows that there are two subscriber sessions running on the same interface. The IP Address field shows that one session is assigned an IPv4 address, and one session is assigned on IPv6 address.

The LS:RI field shows that the subscriber is placed in the correct routing instance and that traffic can be sent and received.

Verifying Dynamic Subscriber Sessions

Purpose

Verify that the dynamic subscriber session is active, and the IPv6 prefix obtained from the ND/RA pool.

Action

From operational mode, enter the `show subscribers detail` command.

```
user@host>show subscribers detail
Type: PPPoE
User Name: dual-stack-v4v6-nas
IP Address: 192.0.2.4
IP Netmask: 255.255.0.0
IPv6 User Prefix: 2001:db8:2010:0:0:6::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741859
Interface type: Dynamic
Dynamic Profile Name: DS-dyn-ipv4v6-ra
```

```
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 81
Session ID: 81
Login Time: 2012-01-17 14:19:41 PST
```

Meaning

The IPv6 User Prefix field shows the prefix that was obtained from the ND/RA pool. The State field shows that the session is active.

Verifying the ND/RA Prefix Pool and Prefix Length

Purpose

Verify the pool used for ND/RA and the prefix length used with the pool

Action

From operational mode, enter the show subscribers extensive command.

```
user@host>show subscribers extensive
Type: PPPoE
User Name: dual-stack-v4v6-nas
IP Address: 192.0.2.4
IP Netmask: 255.255.0.0
IPv6 User Prefix: 2001:db8:2010:0:0:6::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741859
Interface type: Dynamic
Dynamic Profile Name: DS-dyn-ipv4v6-ra
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 81
Session ID: 81
Login Time: 2012-01-17 14:19:41 PST
IPv6 Delegated Address Pool: ndra-2010
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2010:0:0:6::1/64
```

Meaning

Under the PPPoE session, the IPv6 Delegated Address Pool field shows the name of the pool used for ND/RA prefixes. The IPv6 Delegated Network Prefix Length field shows the length of the prefix used to assign the IPv6 address for this subscriber session. The IPv6 Interface Address field shows the IPv6 address assigned to the CPE interface from the ND/RA pool.

Verifying the Status of the PPPoE Logical Interface

Purpose

Display status information about the PPPoE logical interface (pp0).

Action

From operational mode, enter the **show interfaces pp0.logical** command.

```
user@host>show interfaces pp0.1073741859
Logical interface pp0.1073741859 (Index 388) (SNMP ifIndex 674)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 10,
    Session AC name: almach, Remote MAC address: 00:00:5E:00:53:02,
    Underlying interface: ge-3/3/0.1004 (Index 354)
  Bandwidth: 1000mbps
  Input packets : 15
  Output packets: 44
  Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls: Not-configured
  CHAP state: Closed
  PAP state: Success
    Protocol inet, MTU: 65531
      Flags: Sendbroadcast-pkt-to-re
      Addresses, Flags: Is-Primary
        Local: 192.0.2.77
    Protocol inet6, MTU: 65531
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2001:db8:2010:0:0:6::/64, Local: 2001:db8:2010:0:0:6::1
        Local: fe80::2a0:a50f:fc63:a842
```

Meaning

The **Local** field under **Protocol inet** shows the IPv4 address of the pp0 interface. This is the IPv4 address configured for the loopback interface.

The **Destination** field under **Protocol inet6** shows the IPv6 address obtained through ND/RA. This is the value of the *\$junos-ipv6-ndra-prefix* variable configured in the dynamic profile.

The **Local** field under **Protocol inet6** shows the value of the *\$junos-ipv6-address* variable configured for family inet6 in the pp0 configuration of the dynamic profile.

Verifying Router Advertisements

Purpose

Verify that router advertisements are being sent, and router solicits are being received.

Action

From operational mode, enter the **show ipv6 router-advertisement** command.

```
user@host>show ipv6 router-advertisement
Interface: pp0.1073741859
  Advertisements sent: 3, last sent 00:09:53 ago
  Solicits received: 0
  Advertisements received: 0
```

If you have a large number of subscriber interfaces, you can display router advertisements for a specific interface.

```
user@host>show ipv6 router-advertisement interface pp0.1073741859
Interface: pp0.1073741859
  Advertisements sent: 3, last sent 00:10:31 ago
  Solicits received: 0
  Advertisements received: 0
```

Meaning

The display shows the number of advertisements that the router sent, the number of solicits and advertisements that the router received.

SEE ALSO

[Using NDRA to Provide IPv6 WAN Link Addressing Overview | 1250](#)

[Design 3: IPv6 Addressing with NDRA | 1310](#)

[Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | 1353](#)

Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE

IN THIS SECTION

- [Requirements | 1390](#)
- [Overview | 1390](#)
- [Configuration | 1392](#)
- [Verification | 1414](#)

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router
- Junos OS Release 11.4 or later

Overview

IN THIS SECTION

- [Topology | 1391](#)

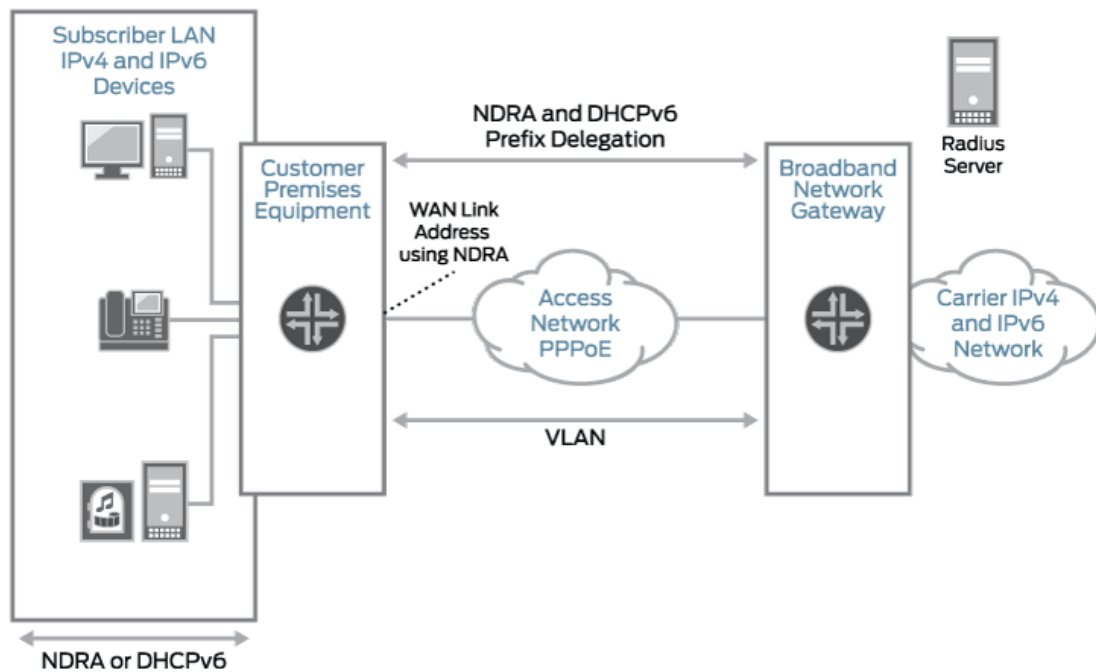
This design uses ND/RA and DHCPv6 prefix delegation in your subscriber access network as follows:

- The access network is PPPoE.
- ND/RA is used to assign a global IPv6 address on the WAN link. The prefixes used in router advertisements come from a local pool that is specified using AAA RADIUS.

- DHCPv6 prefix delegation is used for subscriber LAN addressing. It used a delegated prefix from a local pool that is specified using AAA RADIUS.
- DHCPv4 is used for subscriber LAN addressing.
- DHCPv6 subscriber sessions are layered over an underlying PPPoE subscriber session.

Topology

Figure 25: PPPoE Subscriber Access Network with ND/RA and DHCPv6 Prefix Delegation



8017768

Table 60 on page 1391 describes the configuration components used in this example.

Table 60: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation

Configuration Component	Component Name	Purpose
Dynamic Profiles	DS-dyn-ipv4v6-ndra	Profile that creates a PPPoE logical interface when the subscriber logs in.

Table 60: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation
(Continued)

Configuration Component	Component Name	Purpose
Interfaces	ge-3/3/0	Underlying Ethernet interface.
	lo0	Loopback interface for use in the access network. The loopback interface is automatically used for unnumbered interfaces.
Address-Assignment Pools	default-ipv4-pool-2	Pool that provides IPv4 addresses for the subscriber LAN.
	ndra-2010	Pool that provides IPv6 prefixes used in router advertisements. These prefixes are used to create a global IPv6 address that is assigned to the CPE WAN link.
	dhcpv6-pd-pool	Pool that provides a pool of prefixes that are delegated to the CPE and are used for assigning IPv6 global addresses on the subscriber LAN.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1393](#)
- [Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | 1396](#)
- [Configuring a Dynamic Profile for the PPPoE Logical Interface | 1397](#)
- [Configuring a Loopback Interface | 1401](#)
- [Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces | 1403](#)
- [Specifying the BNG IP Address | 1405](#)
- [Configuring RADIUS Server Access | 1406](#)
- [Configuring RADIUS Server Access Profile | 1408](#)
- [Specifying the RADIUS Server Access Profile to Use | 1410](#)
- [Configuring Local Address-Assignment Pools | 1411](#)

- [Specifying the Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | 1413](#)

CLI Quick Configuration

The following is the complete configuration for this example:

```
dynamic-profiles {
  DS-dyn-ipv4v6-ra {
    interfaces {
      pp0 {
        unit "$junos-interface-unit" {
          ppp-options {
            chap;
            pap;
          }
          pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
          }
          keepalives interval 30;
          family inet {
            unnumbered-address lo0.0;
          }
          family inet6 {
            address $junos-ipv6-address;
          }
        }
      }
    }
  }
  protocols {
    router-advertisement {
      interface "$junos-interface-name" {
        prefix $junos-ipv6-ndra-prefix;
      }
    }
  }
}
system {
```



```

services {
    dhcp-local-server {
        dhcpv6 {
            overrides {
                delegated-pool dhcpv6-pd-pool;
            }
            group DHCPv6-over-pppoe {
                interface pp0.0;
            }
        }
    }
}

access-profile Access-Profile;
interfaces {
    ge-3/3/0 {
        unit 1109 {
            description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
            encapsulation ppp-over-ether;
            vlan-id 1109;
            pppoe-underlying-options {
                duplicate-protection;
                dynamic-profile DS-dyn-ipv4v6-ra;
            }
        }
    }
    lo0 {
        description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
        unit 0 {
            family inet {
                address 192.0.2.77/32 {
                    primary;
                }
            }
            family inet6 {
                address 2001:db8:2030:0:0::1/64 {
                    primary;
                }
            }
        }
    }
}

routing-options {

```

```

    router-id 203.0.113.0;
}
access {
    radius-server {
        203.0.113.99 {
            secret "$ABC123$ABC123ABC123"; ## SECRET-DATA
            timeout 45;
            retry 4;
            source-address 203.0.113.1;
        }
    }
}
profile Access-Profile {
    authentication-order radius;
    radius {
        authentication-server 203.0.113.99;
        accounting-server 203.0.113.99;
    }
    accounting {
        order [ radius none ];
        update-interval 120;
        statistics volume-time;
    }
}
address-assignment {
    pool default-ipv4-pool-2 {
        family inet {
            network 203.0.113.10/16;
            range r5 {
                low 203.0.113.11;
                high 203.0.113.150;
            }
        }
    }
    pool dhcpv6-pd-pool {
        family inet6 {
            prefix 2001:db8:2040:2000:2000::/48;
            range r1 prefix-length 64;
        }
    }
    pool ndra-2010 {
        family inet6 {
            prefix 2001:db8:2010:0:0:0::/48;
            range L prefix-length 64;
        }
    }
}

```

```

    }
  }
}
address-protection;
}

```

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE

CLI Quick Configuration

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

```

edit system services dhcp-local-server dhcpv6
edit group DHCPv6-over-pppoe
set interface pp0.0

```

Step-by-Step Procedure

To layer DHCPv6 above the PPPoE IPv6 family (inet6), associate DHCPv6 with the PPPoE interfaces by adding the PPPoE interfaces to the DHCPv6 local server configuration. Because this example uses a dynamic PPPoE interface, we are using the pp0.0 (PPPoE) logical interface as a wildcard to indicate that a DHCPv6 binding can be made on top of a PPPoE interface.

To configure a DHCPv6 local server:

1. Access the DHCPv6 local server configuration.

```

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

```

2. Create a group for dynamic PPPoE interfaces and assign a name.

The group feature groups a set of interfaces and then applies a common DHCP configuration to the named interface group.

```

[edit system services dhcp-local-server dhcpv6]
user@host# edit group DHCPv6-over-pppoe

```

3. Add an interface for dynamic PPPoE logical interfaces.

```
[edit system services dhcp-local-server dhcpv6 group DHCPv6-over-pppoe]
user@host# set interface pp0.0
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit]
user@host# show
system {
  services {
    dhcp-local-server {
      dhcpv6 {
        group DHCPv6-over-pppoe {
          interface pp0.0;
        }
      }
    }
  }
}
```

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

Configuring a Dynamic Profile for the PPPoE Logical Interface

CLI Quick Configuration

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

```
edit dynamic-profiles DS-dyn-ipv4v6-ra
edit interfaces pp0 unit $junos-interface-unit
set family inet unnumbered-address lo0.0
set family inet6 address $junos-ipv6-address
set pppoe-options underlying-interface "$junos-underlying-interface"
set pppoe-options server
```

```

set ppp-options pap
set ppp-options chap
set keepalives interval 30
up 3
edit protocols router-advertisement
edit interface $junos-interface-name
set prefix $junos-ipv6-ndra-prefix

```

Step-by-Step Procedure

Create a dynamic profile for the PPPoE logical interface. This dynamic profile supports both IPv4 and IPv6 sessions on the same logical interface.

To configure the dynamic profile:

1. Create and name the dynamic profile.

```

[edit]
user@host# edit dynamic-profiles DS-dyn-ipv4v6-ra

```

2. Configure a PPPoE logical interface (pp0) that is used to create logical PPPoE interfaces for the IPv4 and IPv6 subscribers.

```

[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# edit interfaces pp0

```

3. Specify \$junos-interface-unit as the predefined variable to represent the logical unit number for the pp0 interface. The variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

```

[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0]
user@host# edit unit $junos-interface-unit

```

4. Specify \$junos-underlying-interface as the predefined variable to represent the name of the underlying Ethernet interface on which the router creates the dynamic PPPoE logical interface. The variable is

dynamically replaced with the actual name of the underlying interface supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

5. Configure the router to act as a PPPoE server when a PPPoE logical interface is dynamically created.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

6. Configure the IPv4 family for the pp0 interface. Specify the unnumbered address to dynamically create loopback interfaces.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet unnumbered-address lo0.0
```

7. Configure the IPv6 family for the pp0 interface. Because the example uses router advertisement, assign the predefined variable \$junos-ipv6-address.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 unnumbered-address $junos-ipv6-address
```

8. Configure one or more PPP authentication protocols for the pp0 interface.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
user@host# set ppp-options pap
```

9. Enable keepalives and set an interval for keepalives. We recommend an interval of 30 seconds.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra interfaces pp0 unit "$junos-interface-unit"]
user@host# set keepalives interval 30
```

10. Access the router advertisement configuration.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# edit protocols router-advertisement
```

11. Specify the interface on which the ND/RA configuration is applied.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra protocols router-advertisement]
user@host# edit interface $junos-interface-name
```

12. Specify a prefix value contained in router advertisement messages sent to the CPE on interfaces created with this dynamic profile. If you specify the `$junos-ipv6-ndra-prefix` predefined variable, the actual value is obtained from a local pool or through AAA.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra protocols router-advertisement interface "$junos-
interface-name"]
user@host# set prefix $junos-ipv6-ndra-prefix
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit dynamic-profiles DS-dyn-ipv4v6-ra]
user@host# show
interfaces {
  pp0 {
    unit "$junos-interface-unit" {
      ppp-options {
        chap;
        pap;
      }
      pppoe-options {
        underlying-interface "$junos-underlying-interface";
        server;
      }
      keepalives interval 30;
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}
```

```

        family inet6 {
            address $junos-ipv6-address;
        }
    }
}

protocols {
    router-advertisement {
        interface "$junos-interface-name" {
            prefix $junos-ipv6-ndra-prefix;
        }
    }
}

```

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

Configuring a Loopback Interface

CLI Quick Configuration

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

```

edit interfaces lo0 unit 0
set family inet address 192.0.2.77/32 primary
set family inet6 address 2001:db8:2030:0:0::1/64 primary

```

Step-by-Step Procedure

To configure a loopback interface:

1. Create the loopback interface and specify a unit number.

```

[edit]
user@host# edit interfaces lo0 unit 0

```


2. Configure the interface for IPv4.

```
[edit interfaces lo0 unit 0]  
user@host# set family inet address 192.0.2.77/32 primary
```

3. Configure the interface for IPv6.

```
[edit interfaces lo0 unit 0]  
user@host# set family inet6 address 2001:db8:2030:0:0::1/64 primary
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces lo0]  
user@host# show  
unit 0 {  
    family inet {  
        address 192.0.2.77/32 {  
            primary;  
        }  
    }  
    family inet6 {  
        address 2001:db8:2030:0:0::1/64 {  
            primary;  
        }  
    }  
}
```

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

Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces

CLI Quick Configuration

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

```
edit interfaces ge-3/3/0 unit 1109
set description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd"
set encapsulation ppp-over-ether
set vlan-id 1109
set pppoe-underlying-options duplicate-protection
set pppoe-underlying-options dynamic-profile DS-dyn-ipv4v6-ra
```

Step-by-Step Procedure

To configure the underlying Ethernet interface:

1. Specify the name and logical unit number of the static underlying Ethernet interface to which you want to attach the IPv4 and IPv6 dynamic profile.

```
[edit]
user@host# edit interfaces ge-3/3/0 unit 1109
```

2. Configure a description for the interface.

```
[edit interfaces ge-3/3/0 unit 1109]
user@host# set description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd"
```

3. Configure PPPoE encapsulation on the underlying interface.

```
[edit interfaces ge-3/3/0 unit 1109]
user@host# set encapsulation ppp-over-ether
```

4. Configure the VLAN Id.

```
[edit interfaces ge-3/3/0 unit 1109]
user@host# set vlan-id 1109
```

5. Attach the dynamic profile to the underlying interface.

```
[edit interfaces ge-3/3/0 unit 1109]
user@host# set pppoe-underlying-options dynamic-profile DS-dyn-ipv4v6-ra
```

6. (Optional) Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on the same VLAN interface.

```
[edit interfaces ge-3/3/0 unit 1109]
user@host# set pppoe-underlying-options duplicate-protection
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit interfaces]
user@host# show
ge-3/3/0 {
  unit 1109 {
    description "dynamic ipv4v6 dual stack, ndra, dhcpv6 pd";
    encapsulation ppp-over-ether;
    vlan-id 1109;
    pppoe-underlying-options {
      duplicate-protection;
      dynamic-profile DS-dyn-ipv4v6-ra;
    }
  }
}
```

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

Specifying the BNG IP Address

CLI Quick Configuration

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

```
edit routing-options
set router-id 203.0.113.0
```



BEST PRACTICE: We strongly recommend that you configure the BNG IP address to avoid unpredictable behavior if the interface address on a loopback interface changes.

Step-by-Step Procedure

To configure the IP address of the BNG:

1. Access the routing-options configuration.

```
[edit]
user@host# edit routing-options
```

2. Specify the IP address or the BNG.

```
[edit routing-options]
user@host# set router-id 203.0.113.0
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit routing-options]
user@host# show
router-id 203.0.113.0;
```

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

Configuring RADIUS Server Access

CLI Quick Configuration

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

```
edit access radius-server 203.0.113.99
set secret "$ABC123$ABC123ABC123"
set timeout 45
set retry 4
set source-address 203.0.113.1
```

Step-by-Step Procedure

To configure RADIUS servers:

1. Create a RADIUS server configuration, and specify the address of the server.

```
[edit]
user@host# edit access radius-server 203.0.113.99
```

2. Configure the required secret (password) for the server. Secrets enclosed in quotation marks can contain spaces.

```
[edit access radius-server 203.0.113.99]
user@host# set secret "$ABC123$ABC123ABC123"
```

3. Configure the source address that the BNG uses when it sends RADIUS requests to the RADIUS server.

```
[edit access radius-server 203.0.113.99]
user@host# set source address 203.0.113.1
```

4. (Optional) Configure the number of times that the router attempts to contact a RADIUS accounting server. You can configure the router to retry from 1 through 16 times. The default setting is 3 retry attempts.

```
[edit access radius-server 203.0.113.99]
user@host# set retry 4
```

5. (Optional) Configure the length of time that the local router or switch waits to receive a response from a RADIUS server. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

```
[edit access radius-server 203.0.113.99]
user@host# set timeout 45
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
radius-server {
  203.0.113.99 {
    secret "$ABC123$ABC123ABC123"; ## SECRET-DATA
    timeout 45;
    retry 4;
    source-address 203.0.113.1;
  }
}
```

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

Configuring RADIUS Server Access Profile

CLI Quick Configuration

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

```
edit access profile Access-Profile
set authentication-order radius
set radius authentication-server 203.0.113.99
set radius accounting-server 203.0.113.99
set accounting order radius
set accounting order none
set accounting update-interval 120
set accounting statistics volume-time
top
set access-profile Access-Profile
```

Step-by-Step Procedure

To configure a RADIUS server access profile:

1. Create a RADIUS server access profile.

```
[edit]
user@host# edit access profile Access-Profile
```

2. Specify the order in which authentication methods are used.

```
[edit access profile Access-Profile]
user@host# set authentication-order radius
```

3. Specify the address of the RADIUS server used for authentication and the server used for accounting.

```
[edit access profile Access-Profile]
user@host# set radius authentication-server 203.0.113.99
user@host# set radius accounting-server 203.0.113.99
```

4. Configure RADIUS accounting values for the access profile.

```
[edit access profile Access-Profile]
user@host# set accounting order [ radius none ]
user@host# set accounting update-interval 120
user@host# set accounting statistics volume-time
```

5. At the top of the configuration hierarchy, enter the following command to enable the access profile.

```
[edit]
user@host# set access-profile Access-Profile
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit access]
user@host# show
profile Access-Profile {
  authentication-order radius;
  radius {
    authentication-server 203.0.113.99;
    accounting-server 203.0.113.99;
  }
  accounting {
    order [ radius none ];
    update-interval 120;
    statistics volume-time;
  }
}
```


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

Specifying the RADIUS Server Access Profile to Use

CLI Quick Configuration

To quickly configure this example, copy the following command and paste it into the CLI at the `[edit]` hierarchy level.

```
set access-profile Access-Profile
```

Step-by-Step Procedure

To specify the RADIUS server access profile to use for authentication:

1. Specify the access profile.

```
[edit]  
user@host# set access-profile Access-Profile
```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit]  
user@host# show  
...  
access-profile Access-Profile;  
...
```

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

Configuring Local Address-Assignment Pools

CLI Quick Configuration

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

```
edit access
set address-assignment pool default-ipv4-pool-2 family inet network 203.0.113.10/16
set address-assignment pool default-ipv4-pool-2 family inet range r5 low 203.0.113.11
set address-assignment pool default-ipv4-pool-2 family inet range r5 high 203.0.113.150
set address-assignment pool dhcpv6-pd-pool family inet6 prefix 2001:db8:2040:2000:2000::/48
set address-assignment pool dhcpv6-pd-pool family inet6 range r1 prefix-length 64
set address-assignment pool ndra-2010 family inet6 prefix 2001:db8:2010:0:0:0::/48
set address-assignment pool ndra-2010 family inet6 range L prefix-length 64
set address-protection
```

Step-by-Step Procedure

Configure three address-assignment pools for DHCPv4, DHCPv6 prefix delegation, and ND/RA.

To configure the address-assignment pools:

1. Configure the address-assignment pool for DHCPv4.

```
[edit]
user@host# edit access address-assignment pool default-ipv4-pool-2
user@host# edit family inet
user@host# set network 203.0.113.10/16
user@host# set range r5 low 203.0.113.11
user@host# set range r5 high 203.0.113.150
```

2. Configure the address-assignment pool for DHCPv6 prefix delegation

```
[edit]
user@host# edit access address-assignment pool dhcpv6-pd-pool
user@host# edit family inet6
```

```

user@host# set prefix 2001:db8:2040:2000:2000::/48
user@host# set range r1 prefix-length 64

```

3. Configure the address-assignment pool for ND/RA.

```

[edit]
user@host# edit access address-assignment pool ndra-2010
user@host# edit family inet6
user@host# set prefix 2001:db8:2010:0:0:0::/48
user@host# set range L prefix-length 64

```

4. (Optional) Enable duplicate prefix protection.

```

[edit access]
user@host# set address-protection

```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```

[edit access]
user@host# show
address-assignment {
  pool default-ipv4-pool-2 {
    family inet {
      network 203.0.113.10/16;
      range r5 {
        low 203.0.113.11;
        high 203.0.113.150;
      }
    }
  }
  pool dhcpv6-pd-pool {
    family inet6 {
      prefix 2001:db8:2040:2000:2000::/48;
      range r1 prefix-length 64;
    }
  }
  pool ndra-2010 {

```

```

        family inet6 {
            prefix 2001:db8:2010:0:0:0::/48;
            range L prefix-length 64;
        }
    }
}
address-protection;

```

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

Specifying the Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation

CLI Quick Configuration

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

```

edit system services dhcp-local-server dhcpv6
set overrides delegated-pool dhcpv6-pd-pool

```

Step-by-Step Procedure

To specify that the `dhcpv6-pd-pool` is used for DHCPv6 prefix delegation:

1. Access the DHCPv6 local server configuration.

```

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

```

2. Specify the address pool that assigns the delegated prefix.

```

[edit system services dhcp-local-server dhcpv6]
user@host# set overrides delegated-pool dhcpv6-pd-pool

```

Results

From configuration mode, confirm your configuration by entering the `show` command.

```
[edit system]
user@host# show
services {
  dhcp-local-server {
    dhcpv6 {
      overrides {
        delegated-pool dhcpv6-pd-pool;
      }
    }
  }
}
```

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

Verification

IN THIS SECTION

- [Verifying Active Subscriber Sessions | 1414](#)
- [Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | 1415](#)
- [Verifying Dynamic Subscriber Sessions | 1416](#)
- [Verifying DHCPv6 Address Pools Used for NDRA and DHCPv6 Prefix Delegation | 1417](#)
- [Verifying DHCPv6 Address Bindings | 1419](#)
- [Verifying Router Advertisements | 1419](#)
- [Verifying the Status of the PPPoE Logical Interface | 1420](#)

Confirm that the configuration is working properly.

Verifying Active Subscriber Sessions

Purpose

Verify active subscriber sessions.

Action

From operational mode, enter the `show subscribers summary` command.

```
user@host>show subscribers summary
Subscribers by State
  Active: 2
  Total: 2

Subscribers by Client Type
  DHCP: 1
  PPPoE: 1
  Total: 2
```

Meaning

The fields under Subscribers by State show the number of active subscribers.

The fields under Subscribers by Client Type show the number of active DHCP and DHCPoE subscriber sessions.

Verifying Both IPv4 and IPv6 Address in Correct Routing Instance

Purpose

Verify that the subscriber has both an IPv4 and IPv6 address and is placed in the correct routing instance.

Action

From operational mode, enter the `show subscribers` command.

```
user@host>show subscribers
Interface      IP Address/VLAN ID  User Name                LS:RI
pp0.1073741864 203.0.113.5         dual-stack-v4v6-pd      default:default
*              2001:db8:2010:0:0:8::/64
pp0.1073741864 2001:db8:2040:2000:2000:5::/64      default:default
```

Meaning

The Interface field shows that there are two subscriber sessions running on the same interface. The IP Address field shows that one session is assigned an IPv4 address, and one session is assigned on IPv6 address.

The LS:RI field shows that the subscriber is placed in the correct routing instance and that traffic can be sent and received.

Verifying Dynamic Subscriber Sessions

Purpose

Verify dynamic PPPoE and DHCPv6 subscriber sessions. In this example configuration the DHCPv6 subscriber session should be layered over the underlying PPPoE subscriber session.

Action

From operational mode, enter the `show subscribers detail` command.

```
user@host>show subscribers detail
Type: PPPoE
User Name: dual-stack-v4v6-pd
IP Address: 203.0.113.5
IP Netmask: 255.255.0.0
IPv6 User Prefix: 2001:db8:2010:0:0:8::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
Interface type: Dynamic
Dynamic Profile Name: DS-dyn-ipv4v6-ra
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 87
Session ID: 87
Login Time: 2012-01-17 14:45:30 PST

Type: DHCP
IPv6 Prefix: 2001:db8:2040:2000:2000:5::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
```

```

Interface type: Static
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 88
Session ID: 88
Underlying Session ID: 87
Login Time: 2012-01-17 14:46:00 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 07 64 11 07 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

Meaning

When a subscriber has logged in and started both an IPv4 and an IPv6 session, the output shows the active underlying PPPoE session and the active DHCPv6 session.

The Session ID field for the PPPoE session is 87. The Underlying Session ID for the DHCP session is 87, which shows that the PPPoE session is the underlying session.

Verifying DHCPv6 Address Pools Used for NDRA and DHCPv6 Prefix Delegation

Purpose

Verify the pool used for ND/RA, the delegated address pool used for DHCPv6 prefix delegation and the length of the IPv6 prefixes that were delegated to the CPE.

Action

From operational mode, enter the show subscribers extensive command.

```

user@host>show subscribers extensive
Type: PPPoE
User Name: dual-stack-v4v6-pd
IP Address: 203.0.113.5
IP Netmask: 255.255.0.0
IPv6 User Prefix: 2001:db8:2010:0:0:8::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
Interface type: Dynamic

```



```

Dynamic Profile Name: DS-dyn-ipv4v6-ra
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 87
Session ID: 87
Login Time: 2012-01-17 14:45:30 PST
IPv6 Delegated Address Pool: dhcpv6-pd-pool
IPv6 Delegated Network Prefix Length: 64
IPv6 Interface Address: 2001:db8:2040:2000:2000::/48

Type: DHCP
IPv6 Prefix: 2001:db8:2040:2000:2000:5::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
Interface type: Static
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 88
Session ID: 88
Underlying Session ID: 87
Login Time: 2012-01-17 14:46:00 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 07 64 11 07 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 Address Pool: dhcpv6-pd-pool
IPv6 Delegated Network Prefix Length: 64

```

Meaning

Under the PPPoE session, the IPv6 Delegated Address Pool fields show the names of the pools used for DHCPv6 prefix delegation and for ND/RA prefixes. The IPv6 Delegated Network Prefix Length field shows the length of the prefix used to assign the IPv6 address for this subscriber session. The IPv6 Interface Address field shows the IPv6 address assigned to the CPE interface from the ND/RA pool.

Under the DHCP session, the IPv6 Delegated Address Pool fields show the name of the pool used for DHCPv6 prefix delegation. The IPv6 Delegated Network Prefix Length fields shows the length of the prefix used in DHCPv6 prefix delegation.

Verifying DHCPv6 Address Bindings

Purpose

Display the address bindings in the client table on the DHCPv6 local server.

Action

From operational mode, enter the **show dhcpv6 server binding** command.

```
user@host>show dhcpv6 server binding
Prefix          Session Id Expires State   Interface  Client DUID
2001:db8:2040:2000:2000:5::/64  88          86189  BOUND  pp0.1073741864
LL0x1-00:07:64:11:07:02
```

If you have many active subscriber sessions, you can display the server binding for a specific interface.

```
user@host>show dhcpv6 server binding interface pp0.1073741864
Prefix          Session Id Expires State   Interface  Client DUID
2001:db8:2040:2000:2000:5::/64  88          86182  BOUND  pp0.1073741864
LL0x1-00:07:64:11:07:02
```

Meaning

The Prefix field shows the DHCPv6 prefix assigned to the subscriber session from the pool used for DHCPv6 prefix delegation.

Verifying Router Advertisements

Purpose

Verify that router advertisements are being sent, and router solicits are being received.

Action

From operational mode, enter the **show ipv6 router-advertisement** command.

```
user@host>show ipv6 router-advertisement
Interface: pp0.1073741864
```

```

Advertisements sent: 3, last sent 00:03:29 ago
Solicits received: 0
Advertisements received: 0

```

If you have a large number of subscriber interfaces, you can display router advertisements for a specific interface.

```

user@host>show ipv6 router-advertisement interface pp0.1073741864
Interface: pp0.1073741864
  Advertisements sent: 3, last sent 00:03:34 ago
  Solicits received: 0
  Advertisements received: 0

```

Meaning

The display shows the number of advertisements that the router sent, the number of solicits and advertisements that the router received.

Verifying the Status of the PPPoE Logical Interface

Purpose

Display status information about the PPPoE logical interface (pp0).

Action

From operational mode, enter the **show interfaces pp0.logical** command.

```

user@host>show interfaces pp0.1073741864
Logical interface pp0.1073741864 (Index 388) (SNMP ifIndex 681)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 10,
    Session AC name: almach, Remote MAC address: 00:00:5E:00:53:02,
    Underlying interface: ge-3/3/0.1109 (Index 367)
  Bandwidth: 1000mbps
  Input packets : 22
  Output packets: 50
  Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
  LCP state: Opened

```

```

NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls: Not-configured
CHAP state: Closed
PAP state: Success
  Protocol inet, MTU: 65531
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Primary
      Local: 192.0.2.77
  Protocol inet6, MTU: 65531
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 2001:db8:2010:0:8::/64, Local: 2001:db8:2010:0:8::1
      Local: fe80::2a0:a50f:fc63:a842

```

Meaning

The **Underlying interface** field shows the underlying Ethernet interface configured in the example.

The **Destination** field under **Protocol inet6** shows the IPv6 address obtained through ND/RA. This is the value of the *\$junos-ipv6-ndra-prefix* variable configured in the dynamic profile.

The **Local** field under **Protocol inet6** shows the value of the *\$junos-ipv6-address* variable configured for family inet6 in the pp0 configuration of the dynamic profile.

SEE ALSO

[Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | 1309](#)

[DHCPv6 Prefix Delegation over PPPoE | 1261](#)

RELATED DOCUMENTATION

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Dual Stack for PPPoE Access Networks Using NDRA | 1360](#)

[Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1351](#)

[IPv6 WAN Link Addressing with NDRA | 1250](#)

[Designs for IPv6 Addressing in a Subscriber Access Network | 1305](#)

[Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 1259](#)

IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services

SUMMARY

Read this topic to know about the packet triggered subscribers feature available in Junos and how to configure it. Packet triggered subscribers feature creates IP demultiplexing interfaces (IP demux IFL) on receiving a data packet from clients with preassigned IP address.

IN THIS SECTION

- [IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1422](#)
- [IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1424](#)

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

IN THIS SECTION

- [Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services | 1423](#)

Packet triggered subscribers feature creates IP demultiplexing interfaces (IP demux IFL) on receiving a data packet from clients with preassigned IP address. On receiving the first packet, the control plane checks the IP address. If the source IP address matches one of the configured IP address ranges, the subscriber is authenticated with authenticating server. On successful authentication, the IP demux IFL is created using the dynamic profile specified in the CLI. The IP demux IFL adds the framed route and demux source for subscriber using the mask passed by the authenticating server. If the mask is not sent by the authenticating server, access and demux routes are installed using the mask specified in the CLI.

For residential IPv4 subscribers, all traffic from single household typically has same source IPv4 address. Hence, for every household only one IP demux IFL with a single IPv4 address is created. For business IPv4 subscribers, multiple IPv4 addresses may be assigned using framed-routes, resulting in one IP demux IFL representing multiple IPv4 addresses. For IPv6, the source address of traffic coming from same household or business is different as each device has a separate IPv6 address. The most optimal representation of a household or business in this case consists of one IP demux IFL with an IPv6 prefix, representing all IPv6 addresses in the household/business.



NOTE: During IP demux IFL creation if the authentication fails, the IP demux IFL is still created but such IP demux IFL cannot forward any traffic. Any received traffic for the associated subscriber is dropped. All such rejected IP demux IFLs remain in configured state and is referred as configured subscribers. Creating IP demux IFL even if the authentication fails will avoid thrashing as subsequent packets will be dropped on the PFE and will not be punted to the RE. All subscribers in 'Configured' state will be periodically removed. Once these subscribers are removed any new packets received from the same source will get punted to the RE.



NOTE: Packet-Triggered Subscriber support requires that the MAC address of the connected device remain unchanged for the duration of the subscriber session. If the MAC address changes for a packet-triggered subscriber after the subscriber has logged in and the session is up, the subscriber will not be able to connect from the new device with the same IP address. You can avoid this by setting a period during which the session is monitored for subscriber activity. Use the `client-idle-timeout` option at the `[edit access profile profile-name session-options]` hierarchy level. When the timeout expires, the subscriber is gracefully logged out. The subscriber can then successfully log in from the second device. See [Configure Subscriber Session Timeout Options](#).

Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services

- Allows subscriber management and dynamic subscriber interface configuration in cases where devices in the home or business already have IPv4/IPv6 addresses assigned via other means, for example, statically assigned, or via a Cable Modem Termination System (CMTS).
- Supports packet triggered subscribers using authentication and service selection by RADIUS server and allows a maximum of 16 IPv4 and 16 IPv6 address ranges per underlying IFL.
- Allows the authenticating server to pass in the dynamic-profile to use. When the authenticating server passes these values, they take precedence over values configured through CLI.
- Provides throttling mechanism to mitigate DoS-like attack and limit the rate of exception packets sent to RE for IP demux authentication and creation. The throttling mechanism uses the existing DDoS mechanism.

See [Demultiplexing Interface Overview](#)

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

IN THIS SECTION

- [Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services | 1425](#)

Packet triggered subscribers feature creates IP demultiplexing interfaces (IP demux IFL) on receiving a data packet from clients with preassigned IP address. On receiving the first packet, the control plane checks the IP address. If the source IP address matches one of the configured IP address ranges, the subscriber is authenticated with authenticating server. On successful authentication, the IP demux IFL is created using the dynamic profile specified in the CLI. The IP demux IFL adds the framed route and demux source for subscriber using the mask passed by the authenticating server. If the mask is not sent by the authenticating server, access and demux routes are installed using the mask specified in the CLI.

For residential IPv4 subscribers, all traffic from single household typically has same source IPv4 address. Hence, for every household only one IP demux IFL with a single IPv4 address is created. For business IPv4 subscribers, multiple IPv4 addresses may be assigned using framed-routes, resulting in one IP demux IFL representing multiple IPv4 addresses. For IPv6, the source address of traffic coming from same household or business is different as each device has a separate IPv6 address. The most optimal representation of a household or business in this case consists of one IP demux IFL with an IPv6 prefix, representing all IPv6 addresses in the household/business.



NOTE: During IP demux IFL creation if the authentication fails, the IP demux IFL is still created but such IP demux IFL cannot forward any traffic. Any received traffic for the associated subscriber is dropped. All such rejected IP demux IFLs remains in configured state and is referred as configured subscribers. Creating IP demux IFL even if the authentication fails will avoid thrashing as subsequent packets will be dropped on the PFE and will not be punted to the RE. All subscribers in 'Configured' state will be periodically removed. Once these subscribers are removed any new packets received from the same source will get punted to the RE.



NOTE: Packet-Triggered Subscriber support requires that the MAC address of the connected device remain unchanged for the duration of the subscriber session. If the MAC address changes for a packet-triggered subscriber after the subscriber has logged in and the session is up, the subscriber will not be able to connect from the new device with the same IP address. You can avoid this by setting a period during which the session is monitored for subscriber activity. Use the `client-idle-timeout` option at the `[edit access`

profile profile-name session-options] hierarchy level. When the timeout expires, the subscriber is gracefully logged out. The subscriber can then successfully log in from the second device. See [Configure Subscriber Session Timeout Options](#).

Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services

- Allows subscriber management and dynamic subscriber interface configuration in cases where devices in the home or business already have IPv4/IPv6 addresses assigned via other means, for example, statically assigned, or via a Cable Modem Termination System (CMTS).
- Supports packet triggered subscribers using authentication and service selection by RADIUS server and allows a maximum of 16 IPv4 and 16 IPv6 address ranges per underlying IFL.
- Allows the authenticating server to pass in the dynamic-profile to use. When the authenticating server passes these values, they take precedence over values configured through CLI.
- Provides throttling mechanism to mitigate DoS-like attack and limit the rate of exception packets sent to RE for IP demux authentication and creation. The throttling mechanism uses the existing DDoS mechanism.

See [Demultiplexing Interface Overview](#)

Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

IN THIS SECTION

- [Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | 1426](#)
- [On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview | 1426](#)
- [On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview | 1429](#)
- [IPCP Negotiation with Optional Peer IP Address | 1431](#)
- [How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled | 1432](#)
- [Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1433](#)
- [Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1433](#)

- [Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1434](#)
- [Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes | 1434](#)
- [Enabling IPv4 Release Control VSA \(26–164\) in RADIUS Messages | 1435](#)

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

In a dual stack over PPP access network scenario, the dual-stack session remains running as long as either the IPv4 or IPv6 session is active. By default, when a subscriber terminates the IPv4 session, the BNG retains the IPv4 address that was allocated by AAA at login. The address is not released while the dual-stack PPP session is running. If the IPv4 session is renegotiated while the session is running, the same IPv4 address is assigned to the subscriber's IPv4 session. This functionality results in inefficient use of IPv4 addresses.

You can conserve IPv4 addresses by configuring the router to release the IPv4 address if a subscriber is no longer using an IPv4 service. This feature provides on-demand IP address allocation or de-allocation after the initial PPP authentication and IPv6 address or prefix allocation.

The on-demand configuration does not take effect when the destination (peer) IP address is statically configured in the inet family of the PPP interface.

On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview

IN THIS SECTION

- [IPv4 Address Negotiation for Static PPP Subscribers | 1426](#)
- [IPv4 Address Release for Static PPP Subscribers | 1428](#)

This topic describes how on-demand IPv4 address allocation and de-allocation works for static dual-stack PPP subscribers.

IPv4 Address Negotiation for Static PPP Subscribers

The process for IPv4 address negotiation for a static inet family and address over a static PPP interface is as follows:

1. PPP Link Control Protocol (LCP) is established and an IPv4 control protocol is successfully negotiated.
2. The broadband network gateway (BNG) receives an Internet Protocol Control Protocol (IPCP) Configure Request with a 0.0.0.0 IPv4 address option from the customer premises equipment (CPE).
3. The BNG sends an IPCP Configure Request with a local IPv4 address option to the CPE.
4. The BNG sends an Access-Request message with the IPv4-Release-Control VSA (26-164) (if configured) to the RADIUS server.
5. The BNG receives an IPCP Configure ACK from the CPE.
6. The BNG receives an Access-Accept message from the RADIUS server.
 - If a Framed-IP-Address attribute is received, the BNG performs a duplicate address check (if configured). If a duplicate address check is completed successfully, PPP continues IPCP negotiation with the CPE. Otherwise, the entire PPP session is brought down by sending an LCP terminate request to the CPE.
 - If a Framed-Pool attribute is received, then the IPv4 address is allocated from the specified local address pool configured in the BNG. If the IP pool is not configured in the BNG and there is no other IP pool available, the BNG sends an LCP Protocol-Reject message to the CPE.
 - If neither a Framed-IP-Address attribute nor a Framed-Pool attribute is received, then the BNG allocates an IPv4 address from one of the configured local address pools. If the BNG cannot allocate an IPv4 address, the BNG sends an LCP Protocol-Reject message to the CPE.
 - If ADFv4 filters are present in the Access-Accept message, they need to be reinstalled for that subscriber in the BNG.
 - If both IPv4 primary and secondary DNS addresses are present in the Access-Accept message, then both need to be updated for that subscriber in the BNG. If either an IPv4 primary DNS address or an IPv4 secondary DNS address is present in the Access-Accept message, then only the corresponding DNS address needs to be updated for that subscriber.

If an IPv4 address is not available, and the BNG receives an Access-Reject message from the RADIUS server, the following occurs:

- If the Access-Reject message includes the IPv4-Release-Control VSA (26-164), the BNG sends an IPCP terminate request to the CPE. The CPE is then allowed to renegotiate IP NCP.
- If the Access-Reject message does not include the IPv4-Release-Control VSA (26-164), the BNG sends an LCP Protocol-Reject message to the CPE. The CPE must renegotiate the LCP link before it is allowed to renegotiate IP NCP.

If the RADIUS Access-Reject message includes the IPCP Terminate-Request field, the text of Reply Message attribute (18) is appended to the information in the Terminate-Request field, and will be shown in PPP data.

If there is no response from the RADIUS server, then IPCP is terminated.

7. The BNG sends an IPCP Configure NACK with the new IPv4 address option to the CPE.
8. The subscriber secure policy service (if present for inet family) is activated.

The BNG sends an immediate Interim-Accounting message (if configured) with the IPv4-Release-Control VSA (26-164) (if configured) and the Framed-IP-Address attribute to the RADIUS server.

9. The BNG receives an IPCP Configure Request with new IPv4 address option from the CPE.
10. The BNG receives an Interim-Accounting response from the RADIUS server.
11. The BNG sends an IPCP Configure ACK to the CPE.

IPv4 Address Release for Static PPP Subscribers

The process for IPv4 address release for static inet family and address over static PPP interface is as follows:

1. The BNG receives an IPCP terminate request from the CPE.
2. The BNG sends an IPCP terminate ACK to the CPE.
3. The following actions occur:
 - The subscriber secure policy service (if present for inet family) is de-instantiated.
 - If an IPv4 address was allocated from local address pool, the address then becomes available.
 - The IPv4 address entry is cleared from the subscriber record.
 - The BNG sends an immediate Interim-Accounting message (if configured) with the IPv4-Release-Control VSA (26-164) (if configured) to the RADIUS server and the Framed-IP-Address attribute is not included.
 - User Session Statistics are retained for the entire PPP session and are not cleared when the IPv4 address is released.
4. The BNG receives an Interim-Accounting response from the RADIUS server.

No action is taken in the BNG whether or not it receives a response from the RADIUS server.

On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview

IN THIS SECTION

- [IPv4 Address Negotiation for Dynamic PPP Subscribers | 1429](#)
- [IPv4 Address Release for Dynamic PPP Subscribers | 1431](#)

This topic describes how on-demand IPv4 address allocation and de-allocation works for dynamic dual-stack PPP subscribers.

IPv4 Address Negotiation for Dynamic PPP Subscribers

The process for IPv4 address negotiation for a dynamic inet family and address over a static PPP interface is as follows:

1. PPP Link Control Protocol (LCP) is established and IPv6 control protocol is successfully negotiated.
2. The broadband network gateway (BNG) receives an Internet protocol Control Protocol (IPCP) Configure Request with a 0.0.0.0 IPv4 address option from the CPE.
3. The BNG sends an Access-Request message with the IPv4-Release-Control VSA (26-164) (if configured) to the RADIUS server.
4. The BNG receives an Access-Accept message from the RADIUS server.
 - If a Framed-IP-Address attribute is received, then a duplicate address check (if configured) is performed on the BNG. If a duplicate address check is completed successfully, then PPP continues IPCP negotiation with the CPE. Otherwise, the entire PPP session is brought down by sending an LCP terminate request to the CPE.
 - If Framed-Pool attribute is received, then the IPv4 address is allocated from the specified local address pool configured in the BNG. If the pool is not configured in the BNG and there is no other IP pool available, then an IPCP protocol reject is sent to the CPE.
 - If neither a Framed-IP-Address attribute nor a Framed-Pool attribute is received, then the BNG allocates an IPv4 address from one of the configured local address pools. If the BNG cannot allocate an IPv4 address, then an IPCP protocol reject is sent to the CPE.
 - If ADFv4 filters are present in the Access-Accept message, then they need to be reinstalled for that subscriber in the BNG.

- If both IPv4 primary and secondary DNS addresses are present in the Access-Accept message, then both of them need to be updated for that subscriber in the BNG. If either an IPv4 primary DNS address or an IPv4 secondary DNS address is present in the Access-Accept message, then only the corresponding DNS address needs to be updated for that subscriber.

If an IPv4 address is not available, and the BNG receives an Access-Reject message from the RADIUS server, the following occurs:

- If the Access-Reject message includes the IPv4-Release-Control VSA (26-164), the BNG sends an IPCP terminate request to the CPE. The CPE is then allowed to renegotiate IP NCP.
- If the Access-Reject message does not include the IPv4-Release-Control VSA (26-164), the BNG sends an LCP Protocol-Reject message to the CPE. The CPE must renegotiate the LCP link before it is allowed to renegotiate IP NCP.

If the RADIUS Access-Reject message includes the IPCP Terminate-Request field, the text of Reply Message attribute #18 is appended to the information in the Terminate-Request field, and will be shown in PPP data.

If an Access-Challenge message is received instead of an Access-Accept, then the IPCP protocol reject is sent to the CPE.

If there is no response from the RADIUS server, then IPCP is terminated.

5. The BNG sends an IPCP Configure NACK with the new IPv4 address option to the CPE.
6. The dynamic inet family and local address are added and all IPv4 (family inet) services for the dynamic client profile are instantiated.

The BNG sends an IPCP Configure Request with a local IPv4 address option to the CPE.

7. The BNG sends an immediate Interim-Accounting message (if configured) with the IPv4-Release-Control VSA (26-164) (if configured) and a Framed-IP-Address attribute to the RADIUS server.
8. All IPv4 services, such as ascend data filters (ADF) and firewall filters, for the dynamic service profile and the lawful intercept service (if present for inet family) are instantiated and the Service Accounting-Start messages (if service accounting is configured and IPv4 service is not part of a multi-family service profile) are sent to the RADIUS server. If service instantiation fails, then IPCP is terminated and an IPv4 address release process is initiated.
9. The BNG receives an IPCP Configure Request with a new IPv4 address option from the CPE.
10. The BNG sends an IPCP Configure ACK to the CPE.
11. The BNG receives a Service Accounting-Start response from the RADIUS server.
12. The BNG receives an Interim-Accounting response from the RADIUS server.

13. The BNG receives an IPCP Configure ACK from the CPE.

IPv4 Address Release for Dynamic PPP Subscribers

The process for IPv4 address release for dynamic inet family and address over static PPP interface is as follows:

1. The BNG receives an IPCP terminate request from the CPE.
2. The BNG sends an IPCP terminate ACK to the CPE.
3. The following actions occur:
 - All IPv4 (family inet) services for the dynamic client profile are de-instantiated and the dynamic inet family and local address are removed.
 - All IPv4 services, such as ascend data filters (ADF) and firewall filters, for a dynamic service profile and the lawful intercept service (if present for inet family) are de-instantiated. The Service Accounting-Stop messages (if service accounting is configured and IPv4 service is not part of a multi-family service profile) is sent to the RADIUS server.
 - If an IPv4 address was allocated from a local address pool, then it is available.
 - The IPv4 address entry is cleared from the subscriber record
4. The BNG sends an immediate Interim-Accounting message (if configured) with the IPv4-Release-Control VSA (26-164) (if configured) to the RADIUS server and the Framed-IP-Address attribute must not be included.

User Session Statistics and service session statistics for multi-family service are retained for the entire PPP session and is not cleared when the IPv4 address is released.

5. The BNG receives an Interim-Accounting response from the RADIUS server.

No action taken in the BNG whether or not it receives a response from the RADIUS server.

IPCP Negotiation with Optional Peer IP Address

During normal operation for an Internet Protocol Control Protocol (IPCP) negotiation, if the Point-to-Point Protocol (PPP) client does not request a specific IP address, the MX Series server sends an IP address obtained from RADIUS or from the local address pool.

Typically, when the CPE negotiates a statically provisioned IP address, the BNG receives a Framed-IP-Address of 255.255.255.255, and optionally a Framed Route, during PPP authorization. The CPE presents the configured IP WAN address in the IP Address option of the IPCP confReq message. The BNG then accepts the peer's proposed address.

In some other cases, however, the subscriber's public IP address is provisioned locally on the CPE but not explicitly negotiated via IPNCP. If the PPP client seeks a specific IP address, on receiving a NAK from the server, it sends a confReq message without specifying the IP address option. In this case, even though the server sends an IPCP confAck message, the server terminates the client because the server requires an IP address from the client.

You can configure the `peer-ip-address-optional` statement to enable the IPCP negotiation to succeed even though the peer does not include the IP address option in an IPCP configuration request for static and dynamic, and terminated and tunneled, Point-to-Point Protocol over Ethernet (PPPoE) subscribers. By default, this statement is disabled. This feature also supports high availability (HA) and unified in-service software upgrade (unified ISSU).

If the client does not include the IP address option in an IPCP configuration request and the IPCP negotiation succeeds by configuring the `peer-ip-address-optional` statement, then the server does not have the client IP address.



NOTE: If the client does include the IP address option in an IPCP configuration request, it does not matter whether the `peer-ip-address-optional` statement is configured because the subscriber is always available and the server has the client IP address.

An IP address from RADIUS or from the local pool is allocated to the client and the route towards this is added on the server even though the client is not assigned with this address. IPCP is successful and the subscriber becomes available. If you want the server to have the route to the client-requested IP address, use the Framed-Route RADIUS attribute or configure static routes. The client adds or configures static routes towards the server for proper forwarding.

SEE ALSO

Dynamic Profiles Overview

How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled

The following describes the behavior of the border network gateway (BNG) during authentication when on-demand IP address allocation is enabled:

- If the RADIUS server returns a Framed-IP-Address attribute, the BNG does not go to the RADIUS server for address allocation on the first Internet Protocol Control Protocol (IPCP) negotiation. It uses the Framed-IP-Address attribute returned in the initial Access-Accept message. Accounting-Start and periodic Interim-Accounting messages include the Framed-IP-Address attribute. Immediate Interim Accounting messages are not sent to RADIUS server. Address allocation is similar to the process described for a static or dynamic subscriber.

When a Framed-IP-Address is returned from the RADIUS server during authentication and the customer premises equipment (CPE) does not negotiate IPCP, the IPv4 address is not released whether or not the on-demand IP address allocation is enabled.

- If the RADIUS server returns a Framed-Pool attribute, the BNG does not go to the RADIUS server for address allocation upon first IPCP negotiation and it allocates an IPv4 address from the specified local address pool. Accounting-Start and periodic Interim-Accounting messages do not include the Framed-IP-Address attribute until IPCP negotiation. Immediate Interim Accounting messages (if configured) are sent to the RADIUS server. Address allocation is similar to the process described for a static or dynamic subscriber.
- If the RADIUS server does not return either the Framed-IP-Address attribute or the Framed-Pool attribute, address allocation is similar to the process described for a static or dynamic subscriber. Because IPCP is the only Network Control Protocol (NCP) active for these subscribers, the entire PPP session is terminated upon an IPCP terminate request and an Accounting-stop message is sent to the RADIUS server. Immediate Interim-Accounting messages to release the IPv4 address are not sent in this case.

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

To configure static on-demand IPv4 address allocation for dual-stack PPP subscribers:

1. Specify the name and logical unit number of the interface.

```
[edit]
user@host# edit interfaces pp0 unit 1000
```

2. Enable on-demand IP address allocation.

```
[edit interfaces pp0 unit 1000]
user@host# set ppp-options on-demand-ip-address
```

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

To configure dynamic on-demand IPv4 address allocation for dual-stack PPP subscribers:

1. Access the dynamic profile.

```
[edit]
user@host# edit dynamic profiles ppp-dyn-ipv4
```


2. Specify the name and logical unit number of the interface.

```
[edit dynamic profiles ppp-dyn-ipv4]  
user@host# edit interfaces ppp unit 1000
```

3. Enable on-demand IP address allocation.

```
[edit dynamic profiles ppp-dyn-ipv4 interfaces ppp0 unit 1000]  
user@host# set ppp-options on-demand-ip-address
```

Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

To configure static on-demand IP address IPv4 address allocation for dual-stack PPP subscribers at the system level:

1. Specify the protocol.

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

2. Specify the ppp-service option.

```
[edit protocols]  
user@host# edit ppp-service
```

3. Enable on-demand IP address allocation.

```
[edit protocols ppp-service]  
user@host# set on-demand-ip-address
```

Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes

To enable the BNG to send an immediate interim accounting message:

1. Create a profile and assign a name to it.

```
[edit access]
user@host# edit profile profile1
```

2. Under accounting, specify the address-change-immediate-update option.

```
[edit access profile profile1]
user@host# edit accounting
user@host# set address-change-immediate-update
```

Enabling IPv4 Release Control VSA (26–164) in RADIUS Messages

When you are using on-demand address allocation for dual-stack PPP subscribers, you can configure the BNG to include 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.

If no IPv4 address is available during negotiation for static or dynamic PPP subscribers, the RADIUS server includes this VSA in the Access-Reject message it sends to the BNG. The consequence is that the BNG sends an IPCP terminate request to the CPE and the CPE can then renegotiate IPCP.

If you have not enabled VSA 26–164 to be sent, then the Access-Reject message does not include the VSA, and the BNG sends an LCP Protocol-Reject message to the CPE. The CPE must renegotiate the LCP link before it is allowed to renegotiate IP NCP.

The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

You can optionally configure a message that is included in the VSA when it is sent to the RADIUS server.

To enable the IPv4-Release-Control VSA (26-164) in RADIUS messages:

1. Create a profile and assign a name to it.

```
[edit access]
user@host# edit profile profile1
```

2. Specify that you want to configure RADIUS.

```
[edit access profile profile1]
user@host# edit radius
```

3. (Optional) Configure the VSA to be sent and optionally specify a text message to be included in the VSA.

```
user@host# set ip-address-change-notify message message
```

RELATED DOCUMENTATION

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

[Dual Stack for PPPoE Access Networks Using DHCP | 1356](#)

[Dual Stack for PPPoE Access Networks Using NDRA | 1360](#)

Dual Stack Subscribers Monitoring and Management

IN THIS SECTION

- [Monitoring Active Subscriber Sessions | 1436](#)
- [Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance | 1437](#)
- [Monitoring Dynamic Subscriber Sessions | 1438](#)
- [Monitoring Address Pools Used for Subscribers | 1440](#)
- [Monitoring Specific Subscriber Sessions | 1442](#)
- [Monitoring the Status of the PPPoE Logical Interface | 1443](#)
- [Monitoring Service Sessions for Subscribers | 1444](#)
- [Monitoring PPP Options Negotiated with the Remote Peer | 1445](#)
- [Monitoring the RADIUS Attribute Used for NDRA | 1447](#)

Monitoring Active Subscriber Sessions

IN THIS SECTION

- [Purpose | 1437](#)

- [Action | 1437](#)
- [Meaning | 1437](#)

Purpose

View a summary of active subscriber sessions.

Action

From operational mode, enter the `show subscribers summary` command.

```
user@host>show subscribers summary
Subscribers by State
  Active: 2
  Total: 2

Subscribers by Client Type
  DHCP: 1
  PPPoE: 1
  Total: 2
```

Meaning

The output under `Subscribers by State` shows the number of active subscriber sessions.

The output under `Subscribers by Client Type` shows the number of active sessions by type. The two subscriber sessions above represent a DHCPv6 subscriber on a PPPoE access network. When DHCPv6 is layered over PPPoE, two separate subscriber sessions are created for a subscriber.

Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance

IN THIS SECTION

- [Purpose | 1438](#)
- [Action | 1438](#)
- [Meaning | 1438](#)

Purpose

Verify that the subscriber has both an IPv4 and an IPv6 address and is placed in the correct routing-instance.

Action

From operational mode, enter the `show subscribers` command.

```
user@host>show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
pp0.1073741825 203.0.113.162      ipv4-v6-subscriber default:default
pp0.1073741825 2001:DB8::1        default:default
```

Meaning

The Interface field shows that there are two subscriber sessions running on the same interface. The IP Address field shows that one session is assigned an IPv4 address, and one session is assigned on IPv6 address.

The LS:RI field shows that the subscriber is placed in the correct routing instance and that traffic can be sent and received.

Monitoring Dynamic Subscriber Sessions

IN THIS SECTION

- [Purpose | 1438](#)
- [Action | 1439](#)
- [Meaning | 1440](#)

Purpose

Display dynamic PPPoE and DHCPv6 subscriber sessions.

Action

From operational mode, enter the show subscribers detail command.

```

user@host>show subscribers detail
Type: PPPoE
User Name: SBRSTATICUSER
IP Address: 203.0.113.162
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: pppoe-subscriber-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-12-08 09:11:41 PST

Type: DHCP
IPv6 Address: 2001:DB8::1
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-12-08 09:12:11 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 01 02 00 00 01
00 06 00 02 00 03 00 03 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

Meaning

If you are using DHCPv6 over a PPPoE access network, the output shows the relationship of the DHCPv6 subscriber session with its underlying PPPoE subscriber session. In the output for the PPPoE session, the Session ID is 2. The output of the DHCP session shows that the Underlying Session ID is 2.

Monitoring Address Pools Used for Subscribers

IN THIS SECTION

- Purpose | 1440
- Action | 1440
- Meaning | 1441

Purpose

Verify the pool used for NDRA, the delegated address pool used for DHCPv6 prefix delegation, and the length of the IPv6 prefixes that were delegated to the CPE.

Action

From operational mode, enter the `show subscribers extensive` command.

```
user@host>show subscribers extensive
Type: PPPoE
User Name: dual-stack-v4v6-pd
IP Address: 203.0.113.5
IP Netmask: 255.255.0.0
IPv6 User Prefix: 2001:db8:2010:0:0:8::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
Interface type: Dynamic
Dynamic Profile Name: DS-dyn-ipv4v6-ra
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 87
Session ID: 87
Login Time: 2012-01-17 14:45:30 PST
```

```

IPv6 Delegated Address Pool: dhcpv6-pd-pool
IPv6 Delegated Address Pool: ndra-2010
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2010:0:0:8::1/64

Type: DHCP
IPv6 Prefix: 2001:db8:2040:2000:2000:5::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741864
Interface type: Static
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 88
Session ID: 88
Underlying Session ID: 87
Login Time: 2012-01-17 14:46:00 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 07 64 11 07 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 Address Pool: dhcpv6-pd-pool
IPv6 Delegated Network Prefix Length: 64
IPv6 Delegated Network Prefix Length: 48

```

Meaning

Under the PPPoE session, the IPv6 Delegated Address Pool fields show the names of the pools used for DHCPv6 prefix delegation and for NDRA prefixes. The IPv6 Delegated Network Prefix Length field shows the length of the prefix used to assign the IPv6 address for this subscriber session. The IPv6 Interface Address field shows the IPv6 address assigned to the CPE interface from the NDRA pool.

Under the DHCP session, the IPv6 Delegated Address Pool fields show the name of the pool used for DHCPv6 prefix delegation. The IPv6 Delegated Network Prefix Length fields shows the length of the prefix used in DHCPv6 prefix delegation.

Monitoring Specific Subscriber Sessions

IN THIS SECTION

- Purpose | 1442
- Action | 1442
- Meaning | 1443

Purpose

Display information about specific subscriber sessions. If you have many subscriber sessions running, you can use this command to display specific sessions.

Action

From operational mode, enter the `show subscribers extensive id` command.

```
user@host>show subscribers extensive id 2
Type: PPPoE
User Name: SBRSTATICUSER
IP Address: 203.0.113.162
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: pppoe-subscriber-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-12-08 09:11:41 PST

user@host> show subscribers extensive id 3
Type: DHCP
IPv6 Address: 2001:DB8::1
Logical System: default
Routing Instance: default
```

```

Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-12-08 09:12:11 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 01 02 00 00 01
00 06 00 02 00 03 00 03 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

Meaning

The output shows details about specific subscriber sessions.

Monitoring the Status of the PPPoE Logical Interface

IN THIS SECTION

- Purpose | 1443
- Action | 1443
- Meaning | 1444

Purpose

Display status information about the PPPoE logical interface.

Action

```

user@host> show interfaces pp0.1073741888
Logical interface pp0.1073741888 (Index 123) (SNMP ifIndex 707)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
PPPoE:
  State: SessionUp, Session ID: 16,
  Session AC name: centaurus, Remote MAC address: 00:00:5E:00:53:02,
  Underlying interface: ge-1/0/0.1104 (Index 95)

```

```

Input packets : 8
Output packets: 51816
LCP state: Opened
NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Success
Protocol inet, MTU: 1500
  Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Is-Primary
    Local: 192.0.2.77
Protocol inet6, MTU: 1500
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 2001:DB8:0:21::/64, Local: 2001:DB8:0:21::1
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::2a0:a50f:fc61:6d0

```

Meaning

Displays session information about the ppp0 interface.

Monitoring Service Sessions for Subscribers

IN THIS SECTION

- Purpose | 1444
- Action | 1444
- Meaning | 1445

Purpose

Display a details about dual-stack subscriber session.

Action

```

user@host> show subscribers interface pp0.1073741888 extensive
Type: PPPoE
User Name: dual-stack-v4v6-2svc-good

```

```
IP Address: 203.0.113.140
Logical System: default
Routing Instance: default
Interface: pp0.1073741888
Interface type: Dynamic
Dynamic Profile Name: DS-dyn-ipv4v6-ra
MAC Address: 00:00:5E:00:53:02
State: Active
Radius Accounting ID: 155
Session ID: 155
Login Time: 2011-01-30 20:36:53 PST
Service Sessions: 2
```

```
Service Session ID: 174
Service Session Name: l3-v4-service
State: Active
IPv4 Input Filter Name: upstrm-filter-ge-1/0/0.1104-in
IPv4 Output Filter Name: dwnstrm-filter-ge-1/0/0.1104-out
```

```
Service Session ID: 175
Service Session Name: l3-v6-service
State: Active
IPv6 Input Filter Name: v6-up-filter-ge-1/0/0.1104-in
IPv6 Output Filter Name: v6-dn-filter-ge-1/0/0.1104-out
```

Meaning

The highlighted output includes details about a subscriber's service sessions.

Monitoring PPP Options Negotiated with the Remote Peer

IN THIS SECTION

- Purpose | 1446
- Action | 1446

Purpose

Display the PPP options that were negotiated with the CPE. You can also view the IPv4 address that was negotiated with the remote peer. This address matches the address returned from AAA. You can also see this address by using the `show subscribers` command.

Note that this is the only command that will provide the details about the negotiated interface IDs.

Action

```
user@host> show ppp interface pp0.1073741888 extensive
Session pp0.1073741888, Type: PPP, Phase: Network
  LCP
    State: Opened
    Last started: 2011-01-30 20:36:53 PST
    Last completed: 2011-01-30 20:36:53 PST
    Negotiated options:
      Authentication protocol: pap, Magic number: 1174596353, MRU: 1492
  Authentication: PAP
    State: Grant
    Last started: 2011-01-30 20:36:53 PST
    Last completed: 2011-01-30 20:36:53 PST
  IPCP
    State: Opened
    Last started: 2011-01-30 20:36:54 PST
    Last completed: 2011-01-30 20:36:54 PST
    Negotiated options:
      Local address: 192.0.2.77, Remote address: 192.0.2.140
  IPV6CP
    State: Opened
    Last started: 2011-01-30 20:36:54 PST
    Last completed: 2011-01-30 20:36:54 PST
    Negotiated options:
      Local interface identifier: 2a0:a50f:fc61:6d0,
      Remote interface identifier: 200:64ff:fe01:602
```

Monitoring the RADIUS Attribute Used for NDRA

IN THIS SECTION

- [Purpose | 1447](#)
- [Action | 1447](#)

Purpose

Display the RADIUS attribute used for IPv6 NDRA.

Action

To display the RADIUS attribute used for IPv6 Neighbor Discovery router advertisements:

```
host1#show aaa ipv6-nd-ra-prefix
```

```
IPv6 ND RA Prefix      : IPv6-NdRa-Prefix (Juniper VSA)
```

RELATED DOCUMENTATION

[Dual-Stack Access Models in a DHCP Network | 1312](#)

[Dual-Stack Access Models in a PPPoE Network | 1324](#)

4

PART

DHCPv6 for Subscriber Management

- DHCPv6 Local Server | **1449**
 - DHCPv6 Relay Agent | **1455**
 - DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | **1462**
 - DHCPv6 Monitoring and Management | **1464**
-

DHCPv6 Local Server

IN THIS SECTION

- [DHCPv6 Local Server Overview | 1449](#)
- [Configuring DHCPv6 Rapid Commit \(MX Series, EX Series\) | 1450](#)
- [Preventing Binding of Clients That Do Not Support Reconfigure Messages | 1451](#)
- [Configuring the DUID Type Supported by DHCPv6 Servers | 1452](#)
- [Example: Extended DHCPv6 Local Server Configuration | 1453](#)

DHCPv6 Local Server Overview

The DHCPv6 local server is compatible with the DHCP local server and the DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

The DHCPv6 local server provides many of the same features as the DHCP local server, including:

- Configuration for a specific interface or for a group of interfaces
- Site-specific usernames and passwords
- Numbered Ethernet interfaces
- Statically configured CoS and filters
- AAA directed login
- Use of the IA_NA option to assign a specific address to a client

When a DHCPv6 client logs in, the DHCPv6 local server can optionally use the AAA service framework to interact with the RADIUS server. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters.

The client username, which uniquely identifies a subscriber or a DHCP client, must be present in the configuration in order for DHCPv6 local server to use RADIUS authentication.

You can configure DHCPv6 local server to communicate the following attributes to the AAA service framework and RADIUS at login time:

- Client username

- Client password

Based on the attributes that the DHCPv6 local server provides, RADIUS returns the information listed in [Table 61 on page 1450](#) to configure the client:

Table 61: RADIUS Attributes and VSAs for DHCPv6 Local Server

Attribute Number	Attribute Name	Description
27	Session-Timeout	Lease time, in seconds. If not supplied, the lease does not expire
123	Delegated-IPv6-Prefix	Prefix that is delegated to the client
26-143	Max-Clients-Per-Interface	Maximum number of clients allowed per interface

To configure the extended DHCPv6 local server on the router (or switch), you include the `dhcpv6` statement at the `[edit system services dhcp-local-server]` hierarchy level.

You can also include the `dhcpv6` statement at the following hierarchy levels:

- `[edit logical-systems logical-system-name system services dhcp-local-server]`
- `[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server]`
- `[edit routing-instances routing-instance-name system services dhcp-local-server]`

Configuring DHCPv6 Rapid Commit (MX Series, EX Series)

You can configure the DHCPv6 local server to support the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled, the server recognizes the Rapid Commit option in Solicit messages sent from the DHCPv6 client. (DHCPv6 clients are configured separately to include the DHCPv6 Rapid Commit option in the Solicit messages.) The server and client then use a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-message exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, and is beneficial in environments in which networks are under a heavy load.

You can configure the DHCPv6 local server to support the Rapid Commit option globally, for a specific group, or for a specific interface. By default, rapid commit support is disabled on the DHCPv6 local server.

To configure the DHCPv6 local server to support the DHCPv6 Rapid Commit option:

1. Specify that you want to configure the overrides options:

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

2. Enable rapid commit support:

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set rapid-commit
```

SEE ALSO

[Overriding the Default DHCP Local Server Configuration Settings](#) | 1034

Preventing Binding of Clients That Do Not Support Reconfigure Messages

The DHCPv6 client and server negotiate the use of reconfigure messages. When the client can accept reconfigure messages from the server, then the client includes the Reconfigure Accept option in both solicit and request messages sent to the server.

By default, the DHCPv6 server accepts solicit messages from clients regardless of whether they support reconfiguration. You can specify that the server require clients to accept reconfigure messages. In this case, the DHCPv6 server includes the Reconfigure Accept option in both advertise and reply messages when reconfiguration is configured for the client interface. Solicit messages from nonsupporting clients are discarded and the clients are not allowed to bind.

To configure the DHCPv6 local server to bind only clients that support client-initiated reconfiguration:

- Specify strict reconfiguration.

For all DHCPv6 clients:

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

For only a particular group of DHCPv6 clients:

```
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
user@host# set strict
```

The `show dhcpv6 server statistics` command displays a count of solicit messages that the server has discarded.

SEE ALSO

[Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 1201](#)

[Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | 1198](#)

Configuring the DUID Type Supported by DHCPv6 Servers

Every DHCPv6 client and server has a DHCP unique identifier (DUID). Each DUID is globally unique across all DHCPv6 clients and servers in an administrative domain. Messages between clients and servers can carry the client DUID in the Client-Identifier option and the server DUID in the Server-Identifier option. Clients and servers may require that some message types that include different messages may be accepted or discarded based on whether they include one or both of these DUIDs. A server or client may discard some message types when the DUID option value does not match the server's DUID or the client's DUID, respectively.

The DUIDs facilitate communication between client/server pairs by providing a means for each to determine whether it is the intended recipient of a message and also identifying where to forward a response. For example, a server uses the server DUID received in a message from a client to determine whether the message is intended for it. Then it can compare the client DUID it has received against its database. When it finds a match, the server sends the associated configuration information to the client. The server also uses the client DUID to select clients for an Identity Association.

The server DUID conveyed to the client enables the client to distinguish between servers. To target a single server, it may include that DUID when it sends multicast messages; only the server identified by the DUID responds.

RFC, 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6) defines three types of DUIDs, but we support only the DUID-EN and DUID-LL types:

- **DUID-EN—(Supported)** A device vendor assigns a DUID of this type when the device is manufactured. The value consists of the vendor's IANA enterprise number followed by a unique number. This is the default type.

- **DUID-LL—(Supported)** This type of DUID includes a hardware type code recognized by IANA, followed by the link-layer address of any network interface permanently connected to the device. DUID-LL is supported only for DHCPv6 servers.
- **DUID-LLT—(Not supported).** This type is similar to the DUID-LL type, but additionally includes the time that the DUID is generated relative to a specific date and time.

The DUID type is specified per routing instance.

To configure the router to use the DUID-LL type:

- Specify the type.

```
[edit system services dhcp-local-server dhcpv6]
user@host# set server-duid-type duid-ll
```

Remove this configuration to return to supporting the DUID-EN type.

Example: Extended DHCPv6 Local Server Configuration

This example shows a sample extended DHCPv6 local server configuration. The second part of the example shows a sample RADIUS authentication configuration—authentication must be configured for DHCPv6 local server operations.

```
[edit system services]
dhcp-local-server {
  dhcpv6 {
    authentication {
      password $ABC123;
      username-include {
        user-prefix wallybrown;
        domain-name example.com;
      }
    }
  }
  group group_two {
    authentication {
      password $ABC123$ABC123;
      username-include {
        user-prefix south5;
        domain-name example.com;
      }
    }
  }
}
```

```

    }
    interface ge-1/0/3.0;
  }
}

```

The following is a sample RADIUS authentication configuration.

```

[edit access]
radius-server {
  192.168.1.250 {
    port 1812;
    secret $ABC123;
  }
}
profile isp-bos-metro-fiber-basic {
  accounting-order radius;
  authentication-order radius;
  radius {
    authentication-server 192.168.1.250;
    accounting-server 192.168.1.250;
  }
  accounting {
    order radius;
    accounting-stop-on-failure;
    accounting-stop-on-access-deny;
    update-interval 10;
    statistics time;
  }
}

```

RELATED DOCUMENTATION

[DHCP Overview | 1015](#)

[DHCPv6 Relay Agent | 1455](#)

DHCPv6 Relay Agent

IN THIS SECTION

- [DHCPv6 Relay Agent Overview | 1455](#)
- [DHCPv6 Relay Agent Options | 1456](#)
- [Configuring DHCPv6 Relay Agent Options | 1457](#)
- [Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets | 1458](#)
- [Inserting DHCPv6 Remote-ID Option \(Option 37\) In DHCPv6 Packets | 1460](#)

DHCPv6 Relay Agent Overview

When a DHCPv6 client logs in, the DHCPv6 relay agent uses the AAA service framework to interact with the RADIUS server to provide authentication and accounting. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters, such as session timeout and the maximum number of clients allowed per interface.



NOTE: The PTX Series Packet Transport Routers do not support authentication for DHCPv6 relay agents.



NOTE: The following DHCPv6 functionalities are not supported on ACX Series routers:

- Subscriber authentication for DHCPv6 relay agents
- DHCP snooping
- DHCPv6 client
- Liveness detection
- Dynamic profiles
- Option 37 support for remote ID insertion

- Bidirectional Forwarding Detection (BFD) for DHCPv6 relay

The DHCPv6 relay agent is compatible with the DHCP local server and the DHCP relay agent, and can be enabled on the same interface as either the DHCP local server or DHCP relay agent.

To configure the DHCPv6 relay agent on the router (or switch), you include the `dhcpv6` statement at the `[edit forwarding-options dhcp-relay]` hierarchy level.

You can also include the `dhcpv6` statement at the following hierarchy levels:

- `[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]`

See ["DHCPv6 Monitoring and Management" on page 1464](#) for commands specific to viewing and clearing DHCPv6 bindings and statistics.

DHCPv6 Relay Agent Options

You can configure DHCPv6 relay agent to include additional information in the client-originated DHCP packets that the relay agent forwards to a DHCPv6 server. This support is equivalent to the option 82 support provided by the DHCPv4 relay agent. The DHCPv6 server uses the additional information in the packets to determine the IPv6 address to assign to the client. The server might also use the information for other purposes; for example, to determine which services to grant the client, or to provide additional security against threats such as address spoofing. The DHCPv6 server sends its reply back to the DHCPv6 relay agent, and the agent removes the option information from the message, and then forwards the packet to the client.

You can configure the DHCPv6 relay agent to include the following options in the packet the relay agent sends to the DHCPv6 server:

- Relay Agent Interface-ID (option 18)—An ASCII string that identifies the interface on which the client DHCPv6 packet is received. This is the equivalent of the DHCPv4 relay agent option 82 Agent Circuit ID suboption (suboption 1).
- Relay Agent Remote-ID (option 37)—An ASCII string assigned by the DHCPv6 relay agent that securely identifies the client. This is the equivalent of the DHCPv4 relay agent option 82 Agent Remote ID suboption (suboption 2).

Configuring DHCPv6 Relay Agent Options

You can configure DHCPv6 relay agent to insert optional information in the DHCPv6 packets that the relay receives from clients and forwards to a DHCPv6 server. To configure the optional information, you specify the type of information you want to include in the packets. You use the `relay-agent-interface-id` statement to include Relay Agent Interface-ID (option 18) in the packets, or the `relay-agent-remote-id` statement to include Relay Agent Remote-ID (option 37).

When you enable the DHCPv6 options support, you can optionally configure DHCPv6 relay agent to include a prefix or the interface description as part of the option information. For dual-stack environments, you can also specify that the DHCPv6 relay agent use the DHCPv4 option 82 information to populate DHCPv6 option 18 or option 37.

To enable insertion of DHCPv6 options:

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

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

2. Configure DHCPv6 relay agent to insert the Relay Agent Interface-ID option, the Relay Agent Remote-ID option, or both.

- To insert Relay Agent Interface-ID (option 18):

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id
```

- To insert Relay Agent Remote-ID (option 37):

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-remote-id
```

3. (Optional) Specify additional information that you want to include in option 18 or option 37. The `relay-agent-interface-id` and `relay-agent-remote-id` statements both support inclusion of a prefix, interface description, or the DHCPv4 option 82 information. For example:

- To prepend prefix information—This example prepends a prefix that consists of the hostname and logical system name to option 18. You use the `relay-agent-remote-id` statement to add the prefix to option 37.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id prefix host-name logical-system-name
```

- To include the textual interface description—This example uses the description for the device interface instead of the interface identifier in option 18. You use the `relay-agent-remote-id` statement to add the interface description to option 37.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-interface-id use-interface-description device
```

- To use the DHCPv4 option-82 value—This example uses the DHCPv4 option-82 (suboption 2) value for the DHCPv6 option 37 value. You use the `relay-agent-interface-id` statement to use DHCPv4 option 82 (suboption 1) in DHCPv6 option 18.

This example also includes the optional `strict` keyword to specify that the router drops Solicit packets if the packets do not include an option 82 value. If you do not include the `strict` keyword, the router sends the RELAY-FORW message without adding option 37. The `strict` keyword is not supported for the *relay-agent-interface-id* statement.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# set relay-agent-remote-id use-option-82 strict
```

SEE ALSO

[Using DHCP Relay Agent Option 82 Information | 1078](#)

[Using DHCP Relay Agent Option 82 Information | 1078](#)

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

You can configure DHCPv6 relay agent to insert the DHCPv6 Interface-ID (option 18) in the packets that the relay sends to a DHCPv6 server. You can configure the option 18 support at either the DHCPv6 global or group level.

When you configure option 18 support, you can optionally include the following additional information:

- **Prefix**—Specify the `prefix` option to add a prefix to the interface identifier. The prefix can be any combination of hostname, logical system name, and routing instance name.
- **Interface description**—Specify the `use-interface-description` option to include the textual interface description instead of the interface identifier. You can include either the device interface description or the logical interface description.
- **Option 82 Agent Circuit ID suboption (suboption 1)**—Specify the `use-option-82` option to include the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1). This configuration is useful in a dual-stack environment, which has both DHCPv4 and DHCPv6 subscribers that reside over the same underlying logical interface. The router checks for the option 82 suboption 1 value and inserts it into the outgoing packets. If no DHCPv4 binding exists or if the binding does not have an option 82 suboption 1 value, the router sends the packets without adding an option 18.



NOTE: If you specify one of the optional configurations, and the specified information does not exist (for example, there is no interface description), DHCPv6 relay ignores the optional configuration and inserts the default interface identifier in the packets.

To insert the DHCPv6 Interface-ID option (option 18) in DHCPv6 packets:

1. Configure the DHCPv6 relay to include option 18.

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

2. (Optional) Specify the prefix to include in option 18.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set prefix prefix
```

3. (Optional) Specify that option 18 include the textual description of the interface. You can specify either the logical interface description or the device interface description.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-interface-description (logical | device)
```

4. (Optional) Specify that option 18 use the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1) value.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-option-82
```

SEE ALSO

[Using DHCP Relay Agent Option 82 Information | 1078](#)

[Using DHCP Relay Agent Option 82 Information | 1078](#)

Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets

Starting in Junos OS Release 14.1, you can configure DHCPv6 relay agent to insert DHCPv6 Remote-ID (option 37) in the packets that the relay sends to a DHCPv6 server. You can configure option 37 support at either the DHCPv6 global or group level.

When you configure option 37 support, you can optionally include the following information:

- **Prefix**—Specify the `prefix` option to add a prefix to the interface identifier. The prefix can be any combination of hostname, logical system name, and routing instance name.
- **Interface description**—Specify the `use-interface-description` option to include the textual interface description instead of the interface identifier. You can include either the device interface description or the logical interface description.
- **Option 82 Agent Remote-ID suboption (suboption 2)**—Specify the `use-option-82` option to use the value of the DHCPv4 option 82 Remote-ID suboption (suboption 2). This configuration is useful in a dual-stack environment, which has both DHCPv4 and DHCPv6 subscribers that reside over the same underlying logical interface. The router checks for the option 82 suboption 2 value and inserts it into the outgoing packets.



NOTE: If you specify one of the optional configurations, and the specified information does not exist (for example, there is no interface description), DHCPv6 relay ignores the optional configuration and inserts the default interface identifier in the packets.

To insert the DHCPv6 Remote-ID option (option 37) in DHCPv6 packets:

- 1. Configure the DHCPv6 relay to include option 37.

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

- 2. (Optional) Specify the prefix to include with the option 37 information.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]
user@host# set prefix prefix
```

- 3. (Optional) Specify that option 37 include the textual description of the interface. You can specify either the logical interface description or the device interface description.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]
user@host# set use-interface-description (logical | device)
```

- 4. (Optional) Specify that option 37 use the DHCPv4 option 82 Remote-ID suboption (suboption 2) value.

If no DHCPv4 binding exists, or if the binding does not include an option 82 suboption 2 value, by default the router sends the packets without adding option 37. However, you can use the optional strict keyword to specify that the router drop packets that do not have a suboption 2 value.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]
user@host# set use-option-82 strict
```

SEE ALSO

| [Extracting an Option 82 or Option 37 Substring to Create an Interface Set](#) | 1089

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.1	Starting in Junos OS Release 14.1, you can configure DHCPv6 relay agent to insert DHCPv6 Remote-ID (option 37) in the packets that the relay sends to a DHCPv6 server.

RELATED DOCUMENTATION

[DHCPv6 Local Server | 1449](#)[DHCP Overview | 1015](#)

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking

IN THIS SECTION

- [Benefits of Client MAC address Validation | 1464](#)

Starting in Junos OS Release 18.2R1, a nonconfigurable mechanism exists for DHCPv6 local servers and relay agents to drop packets from a client with an unknown MAC address to prevent a malicious client from hijacking a session. When a DHCPv6 local server or relay agent receives a solicit message from a client to establish a session, it extracts the client MAC address (link-layer address) from the message and adds it to a local table that maps MAC addresses to client IPv6 addresses or prefixes. The server or relay agent uses this table to compare MAC addresses received in subsequent messages from the client to validate whether the client is known; if not, it is assumed to be malicious and the control packet is dropped. Because the packet has failed MAC validation, the Client MAC validation counter is incremented.



NOTE: The assumption here is that the client sending the initial solicit message is benign. In this case, client MAC address validation protects against a malicious client trying to hijack a client session that is already established or in the process of being established. The client MAC address validation does not protect against a malicious client that sends the initial solicit message.

When no relay agent is present; the local server shares a link or access node with the client. In this case, the local server extracts the client MAC address directly from the Layer 2 header of the DHCPv6 control packet and validates the address against the table.

When a relay agent is present, validation is performed by the relay agent. *RFC 6939, Client Link-Layer Address Option in DHCPv6*, enables DHCPv6 relay agents that are connected to the same link as a DHCPv6 client to extract the client MAC address from the Ethernet (Layer 2) header in the received

DHCPv6 control packet. The packet includes the client link-layer address as the source MAC address in its Ethernet header. The relay agent validates the MAC address against the value for this client that is stored in its local table. If the address does not match it drops the packet.

If the address is validated by the relay agent and the packet is not dropped, then the relay agent also includes that MAC address in option 79 (Client Link-Layer Address) in the header of the DHCPv6 relay-forward message that the relay agent sends to the local server. When the DHCPv6 local server receives a relay-forward message from a relay agent, the server automatically examines the message for the presence of option 79. When the option is present, the local server extracts the MAC address and validates it against the value stored in the table for this client. If option 79 is not present, the local server cannot perform the validation.

However, because the relay agent has already validated the address, the local server should not discover any address mismatches.

The following scenarios describe possible relay agent configurations and their implications for server validation:

- A single Lightweight DHCPv6 Relay Agent (LDRA; Layer 2) is connected between the client and the server. If the LDRA did not add option 79 to the header, then the local server extracts the client MAC address directly from the Layer 2 header of the DHCPv6 control packet and validates the address against the table.
- One or more Layer 3 DHCPv6 relay agents are connected between the client and the server. In this case, the server checks for option 79 in the header of the innermost relay-forward message sent by the relay agent. The innermost header is viewed because it is the header modified by the first relay agent reached by the client. Other headers are added by subsequent relay agents in the path. These agents do not add option 79 and they cannot extract the MAC address from the first relay agent's Layer 2 header, because that agent changes the address to its own address, as does each subsequent relay agent.
- A combination of a client-facing Layer 2 (LDRA) relay agent followed by one or more Layer 3 DHCPv6 relay agents is connected between the client and the server. The server checks for option 79 in the innermost header of the relay-forward message sent by the relay agent. If the LDRA did not add option 79 to the header, it is probably not capable of changing the MAC address in the header to its own. Consequently, the server next checks the second innermost header for the option, because the first Layer 3 relay agent may have extracted the MAC address and added option 79 to convey the address.

No configuration is required to enable validation of client MAC addresses. You can view how many control packets have been dropped because of a validation failure by issuing the `show dhcpv6 server statistics` command.

Benefits of Client MAC address Validation

- Client MAC address validation enables you to prevent a DHCPv6 client with an unknown MAC address from hijacking a session established by a known client. Usage of DHCPv6 client MAC addresses is likely to increase as it is convenient for correlating DHCPv4 and DHCPv6 clients in a dual-stack environment.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.2R1	Starting in Junos OS Release 18.2R1, a nonconfigurable mechanism exists for DHCPv6 local servers and relay agents to drop packets from a client with an unknown MAC address to prevent a malicious client from hijacking a session.

RELATED DOCUMENTATION

[DHCPv6 Local Server Overview](#) | 1449

[DHCPv6 Relay Agent Overview](#) | 1455

DHCPv6 Monitoring and Management

IN THIS SECTION

- [Requesting DHCPv6 Local Server to Initiate Reconfiguration of Client Bindings](#) | 1465
- [Verifying and Managing DHCPv6 Local Server Configuration](#) | 1466
- [Verifying and Managing DHCPv6 Relay Configuration](#) | 1467

Requesting DHCPv6 Local Server to Initiate Reconfiguration of Client Bindings

You can request that the DHCPv6 local server initiate reconfiguration of all clients or only specified clients.

To request reconfiguration of all clients:

- Specify the all option.

```
user@host> request dhcpv6 server reconfigure all
```

You can use any of the following methods to request reconfiguration of specific clients:

- Specify the IP address of the DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure 2001:db8:1111:2222::
```

- Specify the client ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
```

- Specify the session ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure 5
```

- Specify an interface; reconfiguration is attempted for all clients on this interface.

```
user@host> request dhcpv6 server reconfigure interface fe-0/0/0.100
```

- Specify a logical system; reconfiguration is attempted for all clients or the specified clients in this logical system.

```
user@host> request dhcpv6 server reconfigure all logical-system ls-bldg5
```


- Specify a routing instance; reconfiguration is attempted for all clients or the specified clients in this routing instance.

```
user@host> request dhcpv6 server reconfigure all routing-instance ri-boston
```

SEE ALSO

| [Dynamic Reconfiguration of Clients From a DHCP Local Server | 1197](#)

Verifying and Managing DHCPv6 Local Server Configuration

IN THIS SECTION

- [Purpose | 1466](#)
- [Action | 1466](#)

Purpose

View or clear information about client address bindings and statistics for the DHCPv6 local server.

Action

- To display the address bindings in the client table on the DHCPv6 local server:

```
user@host> show dhcpv6 server binding
```

- To display DHCPv6 local server statistics:

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

- To clear all DHCPv6 local server statistics:

```
user@host> clear dhcpv6 server binding
```

- To clear all DHCPv6 local server statistics:

```
user@host> clear dhcpv6 server statistics
```

Verifying and Managing DHCPv6 Relay Configuration

IN THIS SECTION

- Purpose | [1467](#)
- Action | [1467](#)

Purpose

View or clear address bindings or statistics for extended DHCPv6 relay agent clients:

Action

- To display the address bindings for extended DHCPv6 relay agent clients:

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

- To display extended DHCPv6 relay agent statistics:

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

- To clear the binding state of DHCPv6 relay agent clients:

```
user@host> clear dhcpv6 relay binding
```

- To clear all extended DHCPv6 relay agent statistics:

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

RELATED DOCUMENTATION

[DHCPv6 Local Server](#) | 1449

[DHCPv6 Relay Agent](#) | 1455

5

PART

Packet Triggered Subscriber Services

- [Packet Triggered Subscriber Services | 1470](#)
-

Packet Triggered Subscriber Services

IN THIS CHAPTER

- [IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services | 1470](#)

IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services

SUMMARY

Read this topic to know about the packet triggered subscribers feature available in Junos and how to configure it. Packet triggered subscribers feature creates IP demultiplexing interfaces (IP demux IFL) on receiving a data packet from clients with preassigned IP address.

IN THIS SECTION

- [IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1470](#)
- [Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1472](#)

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

IN THIS SECTION

- [Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services | 1471](#)

Packet triggered subscribers feature creates IP demultiplexing interfaces (IP demux IFL) on receiving a data packet from clients with preassigned IP address. On receiving the first packet, the control plane checks the IP address. If the source IP address matches one of the configured IP address ranges, the subscriber is authenticated with authenticating server. On successful authentication, the IP demux IFL is created using the dynamic profile specified in the CLI. The IP demux IFL adds the framed route and

demux source for subscriber using the mask passed by the authenticating server. If the mask is not sent by the authenticating server, access and demux routes are installed using the mask specified in the CLI.

For residential IPv4 subscribers, all traffic from single household typically has same source IPv4 address. Hence, for every household only one IP demux IFL with a single IPv4 address is created. For business IPv4 subscribers, multiple IPv4 addresses may be assigned using framed-routes, resulting in one IP demux IFL representing multiple IPv4 addresses. For IPv6, the source address of traffic coming from same household or business is different as each device has a separate IPv6 address. The most optimal representation of a household or business in this case consists of one IP demux IFL with an IPv6 prefix, representing all IPv6 addresses in the household/business.



NOTE: During IP demux IFL creation if the authentication fails, the IP demux IFL is still created but such IP demux IFL cannot forward any traffic. Any received traffic for the associated subscriber is dropped. All such rejected IP demux IFLs remains in configured state and is referred as configured subscribers. Creating IP demux IFL even if the authentication fails will avoid thrashing as subsequent packets will be dropped on the PFE and will not be punted to the RE. All subscribers in 'Configured' state will be periodically removed. Once these subscribers are removed any new packets received from the same source will get punted to the RE.



NOTE: Packet-Triggered Subscriber support requires that the MAC address of the connected device remain unchanged for the duration of the subscriber session. If the MAC address changes for a packet-triggered subscriber after the subscriber has logged in and the session is up, the subscriber will not be able to connect from the new device with the same IP address. You can avoid this by setting a period during which the session is monitored for subscriber activity. Use the `client-idle-timeout` option at the `[edit access profile profile-name session-options]` hierarchy level. When the timeout expires, the subscriber is gracefully logged out. The subscriber can then successfully log in from the second device. See [Configure Subscriber Session Timeout Options](#).

Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services

- Allows subscriber management and dynamic subscriber interface configuration in cases where devices in the home or business already have IPv4/IPv6 addresses assigned via other means, for example, statically assigned, or via a Cable Modem Termination System (CMTS).
- Supports packet triggered subscribers using authentication and service selection by RADIUS server and allows a maximum of 16 IPv4 and 16 IPv6 address ranges per underlying IFL.
- Allows the authenticating server to pass in the dynamic-profile to use. When the authenticating server passes these values, they take precedence over values configured through CLI.

- Provides throttling mechanism to mitigate DoS-like attack and limit the rate of exception packets sent to RE for IP demux authentication and creation. The throttling mechanism uses the existing DDoS mechanism.

See [Demultiplexing Interface Overview](#)

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

You can configure the packet triggered subscribers for demux interfaces for both IPv4 and IPv6 addresses. The packet triggered subscribers feature creates IP demux IFL on receiving a data packet from clients with preassigned IP address. Once the IP demux IFL is created framed route and demux source are added for subscriber using the mask passed by the authenticating server.

To enable the packet triggered subscribers feature, configure the demux options in a dynamic profile. Dynamic profiles enable you to dynamically apply configured values to the dynamic interfaces, making them easier to manage.

Before you begin:

- Configure the dynamic profile.

See *Configuring a Basic Dynamic Profile*.

After you configure the dynamic profile, configure the packet triggered subscribers interfaces beginning with the demux interface:

1. Specify that you want to configure the demux interface.

```
[edit interfaces interface-name unit unit-number]
user@host# edit demux
```

2. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

For IPv4:

```
[edit interfaces interface-name unit unit-number demux]
user@host# edit inet
```

For IPv6:

```
[edit interfaces interface-name unit unit-number]
user@host# edit inet6
```



NOTE: The remaining steps all show family inet, but are the same for either family.

3. Specify the demux address type to be based on the source address.

```
[edit interfaces interface-name unit unit-number inet]
user@host# set address source
```

4. Configure the auto-configure details for the family.

```
[edit interfaces interface-name unit unit-number inet]
user@host# edit auto-configure
```

5. Begin the specific packet-triggered subscriber configuration.

```
[edit interfaces interface-name unit unit-number inet address-source]
user@host# edit address-ranges
```

6. Under address range configure the following:

- a. Dynamic profile includes the details for network address, and the range for the demux interface for the family.

```
[edit interfaces interface-name unit unit-number inet address-source address-range]
user@host# edit dynamic-profile
```

- b. Authentication includes the details for password to be included and the username profiles such as, delimiter, domain name, interface name, authentication server, source address and user prefix for the demux interface for the family.

```
[edit interfaces interface-name unit unit-number inet address-source address-range]
user@host# edit authentication
```


Packet Triggered Recovery Support for Static VLAN Subscribers (MX240, MX304, MX480, MX960, MX2010, MX2020, MX10004 and MX10008)

Support for packet triggered functionality based on the following line cards on MX series devices:

- MX304 and other MX series devices with MPC10 line card (ZT ASIC)
- MX10K-LC9600 line card (YT ASIC)

The packet triggered feature supports static IP assigned subscribers, that includes support for:

- Subscribers with and without VLANs.
- Subscribers with IPv4 and IPv6 packets.
- One IP Demux connection per IPv4 or IPv6 address.
- Packet triggered subscribers using authentication and service selection by RADIUS server and SRC.
- Support CoS at subscriber level.
- Throttling mechanism to mitigate DOS-like attack.
- Removal of IP demux interface when no activity is seen for certain configurable duration.

To enable subscriber management service for packet triggered configuration on an underlying interface, use the `enable force` command under `[edit system services hierarchy]` or use the `set system services subscriber-management enable force` command.

```
[edit system services]
  subscriber-management
  {
    enable {
      force;
    }
  }
```



Address-Assignment Pools for Subscriber Management

- [Address-Assignment Pools for Subscriber Management | 1476](#)
-

Address-Assignment Pools for Subscriber Management

IN THIS CHAPTER

- [Address-Assignment Pools for Subscriber Management | 1476](#)

Address-Assignment Pools for Subscriber Management

IN THIS SECTION

- [Address-Assignment Pools Overview | 1477](#)
- [Address Allocation from Linked Address Pools | 1479](#)
- [Address-Assignment Pool Configuration Overview | 1484](#)
- [Configuring an Address-Assignment Pool Name and Addresses | 1486](#)
- [Configuring a Named Address Range for Dynamic Address Assignment | 1486](#)
- [Preventing Addresses from Being Allocated from an Address Pool | 1487](#)
- [Configuring Address-Assignment Pool Usage Threshold Traps | 1489](#)
- [Configuring Address-Assignment Pool Linking | 1490](#)
- [Configuring Address-Assignment Pool Hold-Down | 1492](#)
- [Configuring DHCP Local Address Pool Rapid Drain | 1493](#)
- [Configuring Static Address Assignment | 1495](#)
- [Configuring Duplicate IPv4 Address Protection for AAA | 1496](#)
- [Example: Configuring an Address-Assignment Pool | 1498](#)

Address-Assignment Pools Overview

IN THIS SECTION

- [Address Assignment Types | 1477](#)
- [Named Address Ranges in Address Assignment Pool | 1477](#)
- [Address Allocation from Linked Address Pools | 1478](#)
- [Address Pool Hold-Down State | 1478](#)
- [Address-Assignment Pool for Neighbor Discovery Router Advertisement | 1478](#)
- [Excluding Specified Address or Address Range | 1478](#)
- [Licensing Requirement | 1479](#)
- [Benefits of Address Assignment Pools | 1479](#)

The address-assignment pool enables you to create centralized IPv4 and IPv6 address pools independent of the client applications that use the pools. The authd process manages the pools and the address allocation, whether the addresses come from local pools or from a RADIUS server.

For example, multiple client applications, such as DHCP, can use the same address-assignment pool to provide addresses for their particular clients. Client applications can acquire addresses for either authenticated or unauthenticated clients. The pool selected for a subscriber, based on the RADIUS server or network matching or other rule, is called the matching pool for the subscriber.

Address Assignment Types

Address-assignment pools support both dynamic and static address assignment. In dynamic address assignment, a client is automatically assigned an address from the address-assignment pool. In static address assignment, which is supported for IPv4 pools only, you reserve an address that is then always used by a particular client. Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

Named Address Ranges in Address Assignment Pool

You can configure named address ranges within an address-assignment pool. A named range is a subset of the overall address range. A client application can use named ranges to manage address assignment based on client-specific criteria. For example, for IPv4 address-assignment pools, you might create a named range that is based on a specific DHCP option 82 value. Then, when a DHCP client request matches the specified option 82 value, an address from the specified range is assigned to the client.

Address Allocation from Linked Address Pools

You can link address-assignment pools together to provide backup pools for address assignment. When no addresses are available in the primary or in the matching address pool, the device automatically proceeds to the linked (secondary) address pool to search for an available address to allocate.

Although the first pool in a chain of linked pools is generally considered the primary pool, a matching pool is not necessarily the first pool in the chain.

Lets use an example on how the search mechanism works. Consider a chain of three pools— A, B, and C. Pool A is the primary pool, Pool B is the matching pool for certain subscribers based on information returned by the RADIUS server. The search for an available address for those subscribers uses the following sequence:

- By default, the matching pool (Pool B) is searched first.
- The search moves to the first pool (Pool A) in the chain if address not found.
- The search proceeds through the chain (Pool C) until an available address is found and allocated, or until the search determines no addresses are free.
- In each pool, all address ranges are fully searched for an address.

You can configure the `linked-pool-aggregation` statement to start searching within a block of addresses in each range in the matching pool and then successively through the linked pools. The search then moves back to the first pool in the chain and searches all addresses in all ranges in each pool through the last pool in the chain.

Address Pool Hold-Down State

You can configure an address-assignment pool in hold-down state. When the address pool is in hold-down state, the pool is no longer available to allocate IP addresses for the subscribers. This configuration gracefully transforms the active pool to an inactive state as the previously allocated addresses are returned to the pool. When the pool is inactive, you can safely perform maintenance on the pool without affecting any active subscribers.

Address-Assignment Pool for Neighbor Discovery Router Advertisement

You can explicitly allocate an address-assignment pool for Neighbor Discovery Router Advertisement (NDRA).

Excluding Specified Address or Address Range

For example, you might want to reserve certain addresses or ranges to be used only for static subscribers. When you configure an address or range to be excluded, and the address or an address

within the range, has already been allocated, that subscriber is logged out, the address is deallocated, and the address is marked for exclusion.

Licensing Requirement

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](#) for details, or contact your Juniper Account Team or Juniper Partner.

Benefits of Address Assignment Pools

- The address-assignment pool feature supports both subscriber management and DHCP management.
- You can create centralized pools of addresses independent of client applications.
- You can specify blocks of addresses, named ranges, so that a given address pool can be used to supply different addresses for different client applications or for subscribers that match different sets of criteria.
- You can link pools together to ensure that pools are searched for free addresses in a specific manner, contiguously or noncontiguously.
- You can gracefully transition an address pool from active to inactive by specifying that no further addresses are allocated from the pool.

Address Allocation from Linked Address Pools

You can link address-assignment pools together into a chain to provide backup pools for address assignment. The pool selected for a subscriber, based on the RADIUS server or network matching or some other rule, is called the matching or matched pool for the subscriber. The matching pool might not be the first (primary) pool in the chain. When no addresses are available for allocation from the matching or primary address pool, the router or switch automatically proceeds to another address pool to search for an available address to allocate. When the search discovers no available addresses anywhere, the search stops and no address is allocated for the subscriber.

The search behavior determines not only how the search progresses along a chain of linked pools, but which address ranges within each pool are searched. Depending on where the search starts, your configuration, and whether previously allocated addresses have been freed, the search may continue in the next linked address pool in the chain, or move back to the first pool in the chain.

The search for an available address starts in the pool that matches the subscriber. In many cases, the matching pool is also the first pool in the chain. For some subscribers, the matching pool is farther down the chain. For example, you might configure the RADIUS server to specify the second pool of a chain

rather than the first based on some criteria that it matches during authentication. For another example, you might specify different address ranges for different subscriber groups; whether a particular pool matches a subscriber then depends on which pools are configured for the different address ranges.

The following terms are used to explain the details of the search behavior:

- **lowAddress**—The lowest address in a given range within an address pool.
- **highAddress**—The highest address in a given range within an address pool.
- **nextAddress**—The next address after the last address allocated in a given range within an address pool. This is the address expected to be allocated next. This address, as well as the last range used, is saved as a starting point for searches.

For example, suppose Pool A has a single range that includes the following addresses: 192.0.2.1, 192.0.2.2, 192.0.2.3, 192.0.2.4. In this case, 192.0.2.1 is the **lowAddress** and 192.0.2.4 is the **highAddress**. If 192.0.2.2 was the last address allocated from this pool, then **nextAddress** is 192.0.2.3.

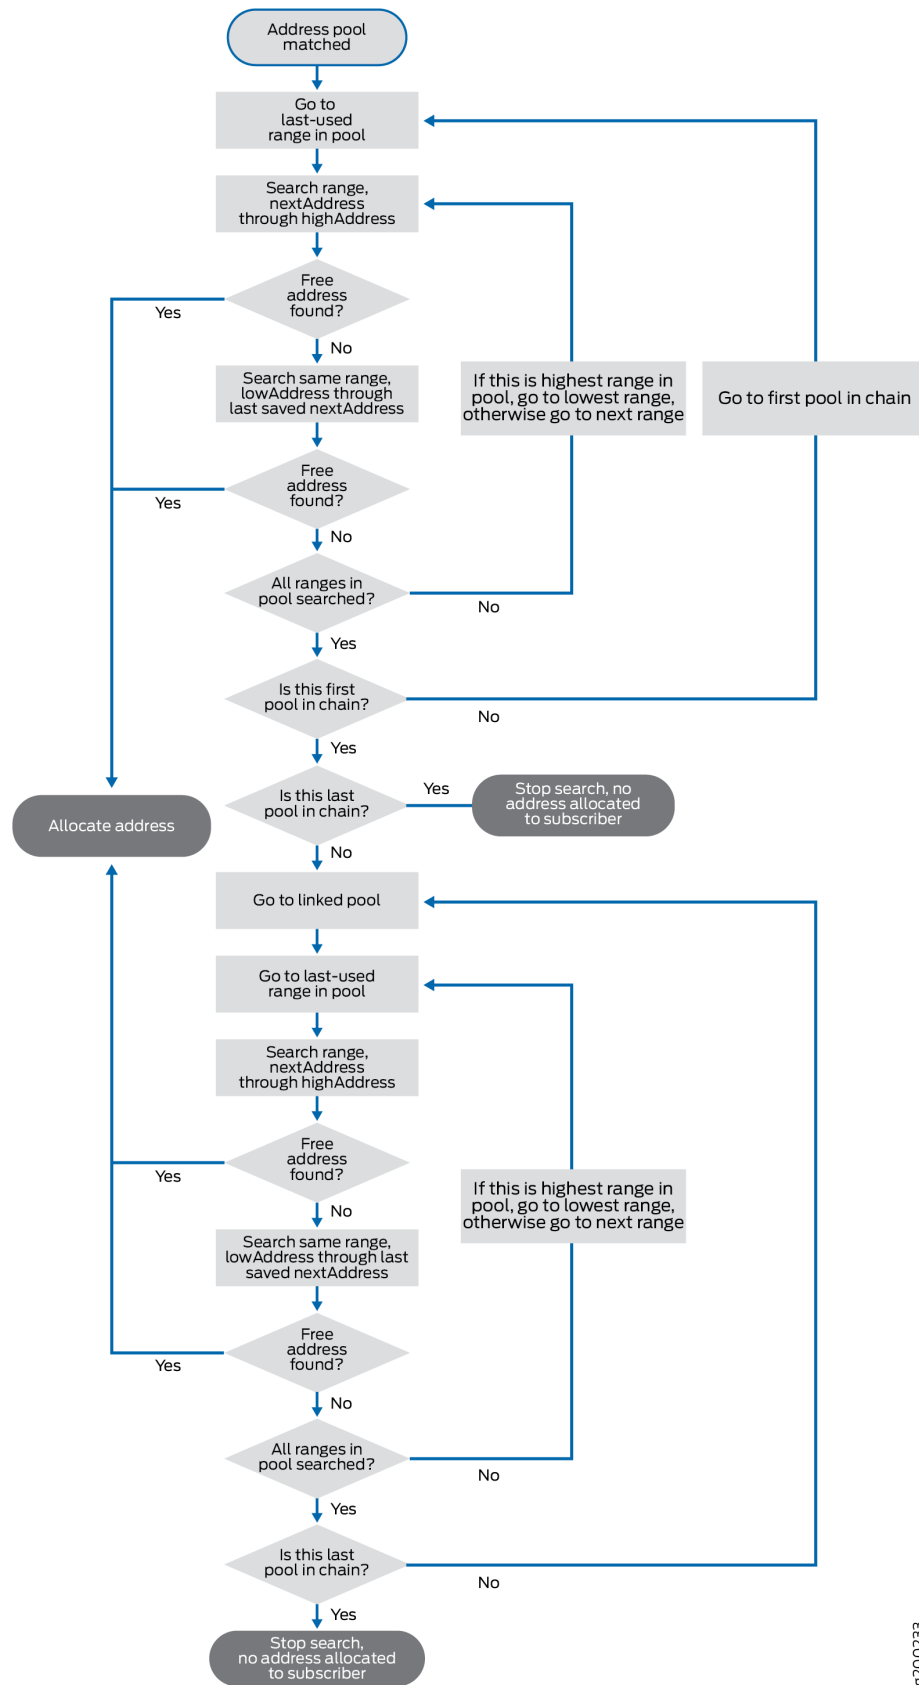
Starting in Junos OS Release 18.1R1, you can configure linked pools to be searched in one of two ways:

- **Contiguous address allocation**—This is the default behavior. All addresses in each range of a pool are searched. The search starts in the matched pool, then moves to the first pool in the chain and, if necessary, continues through each linked pool successively to the last pool in the chain. In each pool, all addresses in all ranges are searched for a free address. This method enables addresses to be assigned contiguously; each pool has to be full before another pool is searched.
- **Noncontiguous (aggregated) address allocation**—Behavior when the `linked-pool-aggregation` statement is configured. Initially, only certain addresses (from **nextAddress** to **highAddress**) are searched in each range of the matched pool. The same search is performed in the linked pool and, if necessary, continues through each linked pool successively to the last pool in the chain.

The search then restarts at the first pool in the chain (not necessarily the matched pool). This time, all addresses in all ranges are searched, in all pools through the end of the chain.

That is the basic functionality, but the details of both searches are fairly complex. [Figure 26 on page 1481](#) shows the default search behavior.

Figure 26: Default Address Assignment from Linked Address Pools



For example, suppose the following conditions exist:

- Linked address pools A, B, C, and D. Pool C is matched.
- Each pool has three address ranges, r1, r2, r3. The last used range was r2 in each pool.

If no free address is found, the search proceeds like this: C > A > B > C > D, then stops.

1. Pool C is searched, nextAddress through highAddress in range r2.
2. Pool C is searched, lowAddress through nextAddress in range r2.
3. Pool C is searched, nextAddress through highAddress in range r3.
4. Pool C is searched, lowAddress through nextAddress in range r3.
5. Pool C is searched, nextAddress through highAddress in range r1.
6. Pool C is searched, lowAddress through nextAddress in range r1.

All ranges and addresses in pool C have been searched, so the search moves to the first pool in the chain, A.


7. Pool A is searched, nextAddress through highAddress in range r2.
8. Pool A is searched, lowAddress through nextAddress in range r2.
9. Pool A is searched, nextAddress through highAddress in range r3.
10. Pool A is searched, lowAddress through nextAddress in range r3.
11. Pool A is searched, nextAddress through highAddress in range r1.
12. Pool A is searched, lowAddress through nextAddress in range r1.

All ranges and addresses in pool A have been searched, so the search moves to the next linked pool in the chain, B.

This process continues until all addresses in all ranges in all pools have been searched. The pool search order is C > A > B > C > D, then stops. Depending on where and whether an address is found, the matching pool might be searched twice. This is true unless the matching pool is the first pool in the chain. For example, if pool A is the matching pool in this set of conditions, then the complete search (assuming no address is found) would be A > B > C > D.

[Figure 27 on page 1483](#) shows the search behavior when you include the linked-pool-aggregation statement.

Figure 27: Aggregated Address Assignment from Linked Address Pools



For example, consider the same conditions exist as for the default example:

- Linked address pools A, B, C, and D. Pool C is matched.
- Each pool has three address ranges, r1, r2, r3. The last used range was r2 in each pool.

If no free address is found, the search proceeds like this: C > D > A > B > C > D, then stops.

1. Pool C is searched, nextAddress through highAddress in range r2.
2. Pool C is searched, nextAddress through highAddress in range r3.
3. Pool C is searched, nextAddress through highAddress in range r1.

All ranges in pool C have been searched from nextAddress to highAddress, so the search moves to the next linked pool in the chain, D.

4. Pool D is searched, nextAddress through highAddress in range r2.
5. Pool D is searched, nextAddress through highAddress in range r3.
6. Pool D is searched, nextAddress through highAddress in range r1.

All ranges in pool D have been searched from nextAddress to highAddress. Pool D is the last pool in the chain, so the search moves to the first pool in the chain, A.

7. Pool A is searched, lowAddress through highAddress in range r2.
8. Pool A is searched, lowAddress through highAddress in range r3.
9. Pool A is searched, lowAddress through highAddress in range r1.

All ranges and addresses in pool A have been searched, so the search moves to the next linked pool in the chain, B.

10. Pool B is searched, lowAddress through highAddress in range r2.
11. Pool B is searched, lowAddress through highAddress in range r3.
12. Pool B is searched, lowAddress through highAddress in range r1.

All ranges and addresses in pool B have been searched, so the search moves to the next linked pool in the chain, C.

This process continues until all addresses in all ranges in all pools have been searched. The pool search order is C > D > A > B > C > D, then stops. Depending on where and whether an address is found, all pools might be searched twice, even when the matching pool is the first pool in the chain. For example, if pool A is the matching pool in this set of conditions, then the complete search (assuming no address is found) would be A > B > C > D > A > B > C > D.

Address-Assignment Pool Configuration Overview

The address-assignment pool feature supports subscriber management functionality by enabling you to create address pools that can be shared by different client applications. An address-assignment pool can support either IPv4 address or IPv6 addresses. You cannot use the same pool for both types of address.



NOTE: Address-assignment pools are completely separate from services PIC-based L2TP LNS address pools, which you create with the `address-pool` statement at the `[edit access]` hierarchy level, and NAT pools, which you create with the `pool` statement at the `[edit services nat]` hierarchy level.

To configure an address-assignment pool:

1. Configure the address-assignment pool name and specify the addresses for the pool.
See ["Configuring an Address-Assignment Pool Name and Addresses" on page 1486](#).
2. (Optional) Configure named ranges (subsets) of addresses.
See ["Configuring a Named Address Range for Dynamic Address Assignment" on page 1486](#).
3. (Optional) Exclude addresses or a range of addresses in address pools from being allocated.
See ["Preventing Addresses from Being Allocated from an Address Pool" on page 1487](#).
4. (Optional) Configure address-assignment pool linking and specify the secondary pool to use when the primary pool is fully allocated.
See ["Configuring Address-Assignment Pool Linking" on page 1490](#).
5. (Optional) Configure address-assignment pool hold-down, so that no additional addresses are allocated from the identified pool. This is also known as passive drain.
See ["Configuring Address-Assignment Pool Hold-Down" on page 1492](#).
6. (Optional) Configure address-assignment pool rapid drain, also known as active drain, to gracefully prevent additional address allocation from the pool and prevent existing clients from renewing leases for addresses from the pool.
See ["Configuring DHCP Local Address Pool Rapid Drain" on page 1493](#).
7. (Optional) Create static address bindings (IPv4 only).
See ["Configuring Static Address Assignment" on page 1495](#).
8. (Optional) Enable duplicate address protection to prevent addresses from being used more than once.
See ["Configuring Duplicate IPv4 Address Protection for AAA" on page 1496](#).
9. (Optional) Configure attributes for DHCP clients.
See ["Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address" on page 1100](#).

SEE ALSO

[DHCP Attributes Overview](#) | 1094

Configuring an Address-Assignment Pool Name and Addresses

To configure an address-assignment pool, you must specify the name of the pool and configure the addresses for the pool.

To configure an IPv4 address-assignment pool:

1. Configure the name of the pool and specify the IPv4 family.

```
[edit access]
user@host# edit address-assignment pool isp_1 family inet
```

2. Configure the network address and the prefix length of the addresses in the pool.

```
[edit access address-assignment pool isp_1 family inet]
user@host# set network 192.168.0.0/16
```

To configure an IPv6 address-assignment pool:

1. Configure the name of the pool and specify the IPv6 family.

```
[edit access]
user@host# edit address-assignment pool isp_2 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 isp_2 family inet6]
user@host# set prefix 2001:db8:2008:2009::/32
```

Configuring a Named Address Range for Dynamic Address Assignment

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.

To create a named range within an IPv4 address-assignment pool:

1. Specify the name of the address-assignment pool and the IPv4 family.

```
[edit access]
user@host# edit address-assignment pool isp_1 family inet
```

2. Configure the name of the range and the lower and upper boundaries of the addresses in the range.

```
[edit access address-assignment pool isp_1 family inet]
user@host# set range southeast low 192.168.102.2 high 192.168.102.254
```

To create a named range within an IPv6 address-assignment pool:

1. Specify the name of the address-assignment pool and the IPv6 family.

```
[edit access]
user@host# edit address-assignment pool isp_2 family inet6
```

2. 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 isp_2 family inet6]
user@host# set range dsl-range low 2001:db8:2008:2010:2011:0100::/64 high
2001:db8:2008:2010:2011:ffff::/64
user@host# set range fiber-east prefix-length 48
```

Preventing Addresses from Being Allocated from an Address Pool

You can exclude specified addresses or address ranges to prevent them from being allocated from an address pool. For example, you might want to reserve certain addresses or ranges to be used only for static subscribers. When you configure an address or range of addresses to be excluded, and the address or an address within the range has already been allocated, that the subscriber allocated that address is logged out, the address is deallocated, and the address is marked for exclusion.

To exclude an address or range of addresses in an address pool from being allocated:

- Specify an individual address.

```
[edit access address-assignment pool-name family (inet | inet6)]
user@host# set excluded-address ip-address
```

- Specify and name a range of consecutive addresses.

```
[edit access address-assignment pool-name family (inet | inet6)]
user@host# set excluded-range name low minimum-value high maximum-value
```

For example, the following partial configuration for an IPv4 address pool defines a range, r1, of addresses to allocate, from 192.168.0.10 through 192.168.128.250. It excludes a single address, 192.168.110.10. It further defines two ranges, exclude1 and exclude2, that specify two sets of consecutive addresses that cannot be allocated from the pool.

```
pool v4-pool {
  family inet {
    network 192.168.0.0/16;
    range r1 {
      low 192.168.0.10;
      high 192.168.128.250;
    }
    excluded-address 192.168.110.10
    excluded-range exclude1 {
      low 192.168.12.0
      high 192.168.12.255
    }
    excluded-range exclude2 {
      low 192.168.98.10
      high 192.168.98.200
    }
  }
}
```

Similarly, the configuration for pool v6-pool defines a range of addresses to allocate and a range of addresses that are excluded from allocation.

```
pool v6-pool {
  family inet6 {
    prefix 2016::/64;
    range r2 {
      low 2016::1;
      high 2016::80:ffff;
    }
    excluded-range exclude3 {
```

```

        low 2016::7c:a
        high 2016::7c:ff
    }
}
}

```

To view information about excluded addresses, you can use either of the following commands:

```
user@host> show network-access address-assignment pool pool-name
```

IP address/prefix	Hardware address	Host/User	Type
192.168.2.1	00:00:5e:00:53:01	user1	DHCP
192.168.2.2	00:00:5e:00:53:02	user2	DHCP
192.168.2.3	00:00:5e:00:53:03	user3	DHCP
192.168.2.4	NA	EXCLUDED	unknown

```
user@host> show network-access aaa statistics address-assignment pool pool-name
```

```
Address-assignment statistics
```

```
...
```

```
Addresses excluded: 1000
```

```
...
```

Configuring Address-Assignment Pool Usage Threshold Traps

You can receive advanced warning that an address pool or linked set of address pools is running short on available addresses by setting usage (utilization) threshold traps. Usage and utilization are used interchangeably to mean the percentage of addresses in an address pool that are currently assigned. An address pool has the following SNMP thresholds associated with it that allow the local address server to signal SNMP traps when certain conditions exist:

- **high-utilization threshold**—When the percentage of addresses assigned from an address pool exceeds this value, a high-utilization SNMP trap is generated. The system sends warning messages as long as this threshold is exceeded.
- **abated-utilization threshold**—When the percentage of addresses assigned from an address pool drops below this value, an abated-utilization trap is generated. The system stops sending the warning messages. Typically, you set the abated-utilization threshold to less than the high-utilization threshold; you cannot set it higher.

The thresholds do not have default values. If you do not configure these thresholds, the system does not send a notification of impending exhaustion as the percentage of addresses assigned approaches 100 percent. The system sends an out-of-address trap only when all addresses in the address pool have been assigned.



NOTE: Starting in Junos OS Release 19.2R1, the out-of-address trap is not sent unless both the high-utilization threshold and the abated-utilization threshold are configured. When the out-of-address trap is sent, an out-of-address syslog message is also sent.



NOTE: You can configure thresholds for all address pools at the [edit access address-assignment] hierarchy level or for only address pools in a specified routing instance at the [edit routing-instance *routing-instance-name*] hierarchy level. The configurations below show only the [edit access] configuration.

To set the threshold traps:

- Specify the high-utilization threshold for IPv4 or IPv6 address pools.

```
[edit accessaddress-assignment]
user@host# set high-utilization percentage
user@host# set high-utilization-v6 percentage
```

- Specify the abated-utilization threshold for IPv4 or IPv6 address pools.

```
[edit accessaddress-assignment]
user@host# set abated-utilization percentage
user@host# set abated-utilization-v6 percentage
```

In the following example, the high threshold is set to 95% usage and the abated threshold is set to 90% usage for IPv4 address pools. When the number of assigned addresses exceeds 95 percent of the address pool, a high-utilization trap is generated. If all the addresses become assigned from the pool, an out-of-address trap is generated and an out-of-address syslog message is sent. When the number of assigned addresses drops below 90 percent of the address pool, the abated-utilization trap is generated.

```
[edit accessaddress-assignment]
user@host# set high-utilization 95
user@host# set abated-utilization 90
```

Configuring Address-Assignment Pool Linking

Address-assignment pool linking enables you to specify a secondary address pool for the router to use when the matching or primary address-assignment pool is fully allocated. You can create a chain of

multiple linked pools. For example, you can link Pool A to Pool B, and link Pool B to Pool C. You can link any number of pools serially in a chain, but you cannot create multiple links to or from the same pool. For example, you cannot create links from Pool A to both Pool B and Pool C. Similarly, Pool C cannot be linked from both Pool A and Pool B. An additional consideration is that all address pools in a chain must be of the same family type, IPv4 or IPv6.

When the address pool that matches the subscribers has no available addresses, the router automatically switches over to the linked pool and allocates addresses from that pool. The router uses a linked pool only when the matching address-assignment pool is fully allocated.

Starting in Junos OS Release 18.1, the behavior changes for how to find and allocate a free address in a chain of address pools. You can configure linked pools to be searched in one of two ways:

- **Contiguous address allocation**—Default behavior. All addresses in each range of a pool are searched. The search starts in the matched pool, then moves to the first pool in the chain and, if necessary, continues through each linked pool successively to the last pool in the chain. In each pool, all addresses in all ranges are searched for a free address. This method enables addresses to be assigned contiguously; each pool has to be full before another pool is searched.
- **Noncontiguous (aggregated) address allocation**—Behavior when `linked-pool-aggregation` is configured. Initially, only certain addresses (from `nextAddress` to `highAddress`) are searched in each range of the matched pool. The same search is performed in the linked pool, if necessary, and continues through each successive linked pool through the last pool in the chain.

The search then restarts at the first pool in the chain (not necessarily the matched pool). This time, all addresses in all ranges are searched, in all pools through the end of the chain.

Including the `linked-pool-aggregation` statement might be desirable if you configure your RADIUS server to use the IP address alone to identify subscribers. Typically, subscribers are identified by the RADIUS server using the subscriber session ID and other criteria. If you use only the IP address, you might encounter the following issue with the default behavior when the `linked-pool-aggregation` statement is not configured. A subscriber can disconnect and that address can be assigned to the next subscriber. The `Acct-Start` message for the second subscriber might be sent before the `Acct-Stop` message is sent for the disconnected subscriber. When the `Acct-Stop` is received, the new subscriber, identified only by the IP address, can be disconnected.

You can avoid this situation by either including the `linked-pool-aggregation` statement or configuring your RADIUS server to use the subscriber session ID (instead of the IP address) for identification.

Before you begin, configure your address pools. See ["Address-Assignment Pool Configuration Overview" on page 1484](#).

To link an address-assignment pool to a secondary pool:

1. Specify the names of the pools to be linked.

```
[edit access address-assignment pool-name]
user@host# set link secondary-pool-name
```

2. (Optional) Configure searching to allow for noncontiguous address allocation.

```
[edit access]
user@host# set linked-pool-aggregation
```

For example, the following configuration links Pool_A to Pool_B and then links Pool_B to Pool_C.

```
[edit access]
user@host# set address-assignment pool Pool_A link Pool_B
user@host# set address-assignment pool Pool_B link Pool_C
```

Configuring Address-Assignment Pool Hold-Down

The address-assignment pool hold-down feature—also known as passive drain—enables you to gracefully transition an active address pool to an inactive state. When the pool is in the inactive state, you can safely perform maintenance on the pool without affecting any current subscribers (such as adding, changing, or deleting addresses).

When an address-assignment pool is in the hold-down state, no additional addresses are allocated from that pool. However, the hold-down state does not affect any existing subscribers that are using addresses previously assigned from the pool. As the existing subscribers disconnect, their IP addresses are marked as free in the pool, but the addresses are not reallocated due to the pool's hold-down state. Eventually, when all subscribers have disconnected and their addresses are returned to the pool, the pool becomes inactive.

To place an active address-assignment pool in the hold-down state:

1. Specify the name of the address-assignment pool.

```
[edit access]
user@host# edit address-assignment pool isp_1
```

2. Specify that the pool is in the hold-down state so that no additional addresses can be allocated from the pool.

```
[edit access address-assignment pool isp_1]  
user@host# set hold-down
```

Configuring DHCP Local Address Pool Rapid Drain

You can force the DHCP local server to stop allocating addresses from a specific local address pool by configuring the pool to active-drain mode. This mode enables the server to gracefully terminate subscribers that are already using addresses assigned from that pool and transition them to another pool. When a DHCP subscriber attempts to renew (at the T1 renewal time) the lease on an IP address from a pool now configured for active-drain mode, the DHCP local server replies with a NAK to the subscriber's renewal request. This response forces the subscriber to renegotiate a lease. The server then allocates a new IP address from an alternative address pool that is not configured for active drain.

The active-drain mode provides a way to rapidly drain subscribers from an address pool. Consequently, the longer the configured lease time for subscribers, the more useful active-drain mode may be. If you do not configure active-drain mode for a pool, then to stop the allocation of its addresses, you must either configure passive-drain mode or delete the pool.

- Passive-drain mode places the address pool in a hold-down state. No more addresses are allocated from the pool, but subscribers currently using an assigned address from the pool are not affected. Existing subscribers are allowed to age out. When the subscriber disconnects (or is disconnected by an operator) the address is released, but cannot be reassigned. Eventually, all subscribers have released their addresses, and the pool is no longer active. Because leases for active subscribers are renewed on request, passive-drain mode can take much longer than active-drain mode to recover all addresses in the pool.
- Pool deletion disrupts the traffic for each current subscriber using a pool address for as long as it takes for the lease to expire and for the subscriber to renegotiate and obtain a new lease. The server removes all subscribers with an address from the deleted pool. The subscribers attempt to extend the lease but fail because the lease was deleted at the server. When the subscribers subsequently attempt to renegotiate a new lease, it may be granted with an address from a different pool or from RADIUS.

You can delete the active-drain configuration before the address pool is emptied. In this case, lease extensions may be granted for subscribers still having addresses from this pool. This recovery is best effort, because some subscribers are in the process of being logged out by the server when the configuration is deleted. These subscribers cannot be recovered to the pool and must renegotiate a lease. These subscribers might then be assigned an address either from this pool (because it is active again) or from an alternate pool.

If the DHCP client fails to receive notification that the address pool is being drained, it may continue to grant lease extensions to subscribers using this pool. This condition is indicated when the address remains bound to the client beyond the T1 time (up to the T2 time) when it should have been recovered by the pool. In this situation, delete the active-drain configuration, then reconfigure it for the pool to ensure the pool is drained in a timely manner.

In the event of an authd or jdhcpd restart, or of a graceful Routing Engine switchover, pool addresses might still be used by some subscribers for whom a NAK has not been sent to initiate the logout. When the restart or GRES completes, authd sends jdhcpd a notification with a list of subscribers still having addresses from the pool that is configured for active drain. Pool draining can then continue.



NOTE: Starting in Junos OS Release 18.4R1, the method of address allocation determines the subsequent behavior when authd notifies the DHCP process that an address pool is deleted or being drained.

- When addresses are allocated on demand, the family with the address in that pool is logged out immediately when the pool is deleted, or logged out gracefully by the draining process when a DHCP renew or rebind message is received.
- When the addresses are preallocated, the addresses for both families are deleted immediately when the pool is deleted, or deleted gracefully by the draining process when a DHCP renew or rebind message is received.

To configure the DHCP local server to stop allocating addresses from an address pool:

1. Access the address pool configuration.

```
[edit access]
user@host# edit address-assignment pool pool-name
```

2. Specify active drain mode for the pool.

```
[edit access address-assignment pool pool-name]
user@host# set active-drain
```

You can use the `show network-access aaa statistics address-assignment pool pool1` command to confirm that active drain is configured for a pool.

```
user@host> show network-access aaa statistics address-assignment pool pool1
Address assignment statistics
Pool Name: pool1
Out of Memory: 0
```

```

Out of Addresses: 0
Address total: 33009
Addresses in use: 1
Address Usage (percent): 0
Pool drain configured: yes

```



NOTE: The active-drain feature takes precedence over preservation of the prefix address. Address preservation may ensure that the same delegated prefix is assigned to the subscriber based on the access circuit identifier (ACI). When a subscriber with a preserved prefix logs out, the ACI and prefix are stored in the address preservation table. When that subscriber tries to log in again, the address and ACI are looked up in the table.

Active drain mode affects this behavior. When the prefix is currently part of a pool set to active-drain mode, it is removed from the table and is not assigned to the subscriber when the subscriber tries to log in again.

If active drain is cancelled while the client is in the process of logging out, then the prefix and ACI string are preserved in the table. In this case, the prefix can be assigned to that ACI string when the subscriber logs in again. However, if active drain is cancelled after the client has already logged out and the table has been cleared of the prefix/ACI association, then the subscriber at a subsequent login gets a prefix from the pool that is reactivated and the prefix could be different.

SEE ALSO

[Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 1100](#)

[Attributes That Can Be Applied to DHCP Clients | 1095](#)

Configuring Static Address Assignment

You can optionally create a static IPv4 address binding by reserving a specific address for a particular client. The address is removed from the address-assignment pool so that it is not assigned to another client. When you reserve an address, you identify the client host and create a binding between the client MAC address and the assigned IP address. IPv6 address-assignment pools do not support static address binding.

To configure a static binding for an IPv4 address:

1. Specify the name of the IPv4 address-assignment pool containing the IP address you want to reserve for the client.

```
[edit access]
user@host# edit address-assignment pool isp_1 family inet
```

2. Specify the name of the client for the static binding, the client MAC address, and the IP address to reserve for the client. This configuration specifies that the client with MAC address 00:00:5E:00:53:90 is always assigned IP address 192.168.44.12.

```
[edit access address-assignment pool isp_1 family inet]
user@host# set host sva1e6_boston_net hardware-address 00:00:5E:00:53:90 ip-address
192.168.44.12
```

Configuring Duplicate IPv4 Address Protection for AAA

Starting in Junos OS Release 14.1, if you are using AAA to supply IPv4 addresses, you can enable duplicate address protection to prevent addresses from being used more than once. If enabled, the following attributes received from external servers are checked:

- *Framed-IP-Address*
- *Framed-Pool*

The router then takes one of the following actions:

- If an address matches an address in an address pool, the address is taken from the pool, provided it is available.
- If the address is already in use, it is rejected as unavailable and the existing subscriber using the address remains intact.

To configure duplicate address protection:

1. Enter the access configuration.

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

2. Enable duplicate address protection.

```
[edit access]
user@host# set address-protection
```

Starting in Junos OS Release 18.4R1, you can optionally enable the reassignment of an address that is currently in use when address protection is configured by including the `reassign-on-match` option. When configured, the router disconnects the existing subscriber and allows the new subscriber to renegotiate. The effect of this configuration is that the address in use is always reassigned to the new subscriber.

One use case for this override capability occurs when a mobile subscriber is accidentally dropped from the gateway GPRS support node (GGSN), but the GGSN keeps the subscriber's L2TP session up for some period of time. When the customer tries to reconnect through a different node, the session cannot connect because the original session is still up. Address reassignment enables the new session to preempt the existing session, allowing the subscriber to reconnect.



NOTE: The existing subscriber is disconnected only when the address is from a RADIUS-sourced address pool. When the address is from a locally configured address pool, the existing subscriber session remains intact.



BEST PRACTICE: Do not use the `reassign-on-match` option when RADIUS is allocating addresses that are contained in a locally configured address pool because there is a greater chance of address collision. We recommend that you do not overlap RADIUS-sourced addresses with local address pools.

The `reassign-on-match` option works in the following way:

1. A subscriber negotiates access with a given IP address.
2. The router determines whether that address is in use and where it came from.
 - When a subscriber is already logged in with that address, the address is not part of a locally configured pool, and address protection is enabled:
 - The router sends a NAK to the new subscriber, rejecting the request.
 - The router sends a disconnect request to the existing subscriber. The disconnect request includes a termination ID to report the cause of the logout.
 - The new (rejected) subscriber can renegotiate and is assigned the IP address.
 - When a subscriber is already logged in with that address and the address was allocated from a local address pool:

- The router sends a NAK to the new subscriber, rejecting the request.
- The router does not send a disconnect request to the existing subscriber.

When you add `reassign-on-match` to an existing duplicate address protection configuration, it takes effect immediately for the existing subscribers. Similarly, if you remove `reassign-on-match` from a configuration, it takes effect immediately, so that a subsequent request for access with an in-use address does not result in termination of the existing subscriber.

To enable address reassignment:

-

```
[edit access address-protection]
user@host# set reassign-on-match
```

SEE ALSO

[Configuring Duplicate IPv6 Prefix Protection for Router Advertisement](#) | 1363

Example: Configuring an Address-Assignment Pool

IN THIS SECTION

- [Requirements](#) | 1498
- [Overview](#) | 1499
- [Configuration](#) | 1499

Requirements

Create the network access address assignment pool. To expand the pool automatically when it reaches 100% usage, to match the new pool ranges to the original pool size, use a different netmask for each pool.



NOTE: There is a restriction for IP range when we link two address assignment pools. We cannot create 2 address pools with same network prefix. We can create address pools with different network prefix and link them if required.

Overview

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1499](#)

This example shows an address-assignment pool configuration that creates two pools, one for IPv4 *DHCP* clients (isp_1), and a second pool (chi-fiber-ra) that is used for router advertisement.

CLI Quick Configuration

```
[edit access]
address-assignment {
  network-discovery-router-advertisement chi-fiber-ra;
  pool isp_1 {
    family inet {
      network 192.168.0.0/16;
      range southeast {
        low 192.168.102.2 high 192.168.102.254;
      }
      range northeast {
        low 192.168.119.2 high 192.168.119.250;
      }
      host host.example.net {
        hardware-address 00:00:5E:00:53:90;
        ip-address 192.168.44.12;
      }
      dhcp-attributes {
        option-match {
          option-82 {
            circuit-id fiber range northeast;
          }
          option-82 {
            circuit-id cable_net range southeast;
          }
        }
      }
    }
  }
}
```

```

        boot-file boot.client;
        boot-server 192.168.200.100;
        grace-period 3600;
        maximum-lease-time 18000;
        netbios-node-type p-node;
        router 192.168.44.44 192.168.44.45;
    }
}
}
pool chi-fiber-ra {
    family inet6 {
        prefix 2001:db8:2008:2009:2010::/48;
        range fiber3 {
            low 2001:db8:2008:2009:2010::1/64;
            high 2001:db8:2008:2009:2010::5/64;
        }
    }
}
}
```

This example creates an IPv4 address-assignment pool named `isp-1`, which contains two named address ranges, `southeast` and `northeast`. The address-assignment pool also contains a static binding for client host `host.example.net`. The `ISP_1` pool configuration also includes the `dhcp-attributes` statement, indicating that the pool is used for DHCP clients. If the option 82 `circuit-id` entry matches the string `fiber`, then DHCP assigns the client an address from the `northeast` range. If the option 82 `circuit-id` matches the string `cable_net`, DHCP assigns an address from the `southeast` range.

The second address-assignment pool created in this example is `chi-fiber-ra`. The `neighbor-discovery-router-advertisement` statement at the beginning of the syntax specifies that this named address-assignment pool is used for router advertisement. The syntax at the end of the example configures the address-assignment pool named `chi-fiber-ra`.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, the out-of-address trap is not sent unless both the high-utilization threshold and the abated-utilization threshold are configured.
18.4R1	Starting in Junos OS Release 18.4R1, you can optionally enable the reassignment of an address that is currently in use when address protection is configured by including the <code>reassign-on-match</code> option.

18.1R1	Starting in Junos OS Release 18.1R1, search mechanism for an available address proceeds through a chain of linked pools. This behavior enables the DHCP to search addresses contiguously.
18.1R1	Starting in Junos OS Release 18.1R1, you can exclude a specified address or range of consecutive addresses to prevent them from being allocated from an address pool.
18.1R1	Starting in Junos OS Release 18.1, the behavior changes for how to find and allocate a free address in a chain of address pools.
14.1	Starting in Junos OS Release 14.1, if you are using AAA to supply IPv4 addresses, you can enable duplicate address protection to prevent addresses from being used more than once.

7

PART

DNS Addresses for Subscriber Management

- [DNS Addresses for Subscriber Management | 1503](#)
-

DNS Addresses for Subscriber Management

IN THIS CHAPTER

- [DNS Name Server Addresses for Subscriber Management | 1503](#)

DNS Name Server Addresses for Subscriber Management

IN THIS SECTION

- [DNS Name Server Address Overview | 1503](#)
- [Configuring DNS Name Server Addresses for Subscriber Management | 1505](#)
- [Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | 1507](#)
- [DNS Resolver for IPv6 DNS Overview | 1508](#)
- [Configuring a DNS Server Address for IPv6 Hosts | 1508](#)

DNS Name Server Address Overview

IN THIS SECTION

- [Benefits of Local DNS Server Addresses | 1504](#)

When a client attempts to access a domain—for example, `www.example.com`—a request is sent to a Domain Name System (DNS) name server. The name server stores information that correlates domain names with IP addresses; the IP address is used to reach the requested domain. In response to the client request, the name server looks up the IP address for the domain—`192.0.2.10` for `www.example.com`—and returns it to the client.

In your network configuration, you must configure the address of one or more name servers locally on the router or on your RADIUS server. The local configuration supports the following subscriber types:

- DHCPv4 or DHCPv6
- IP over Ethernet (VLAN)
- Terminated PPPoE (IPv4 or IPv6)
- Tunneled PPPoE (IPv4 or IPv6)

You can configure the name server addresses globally (per routing instance), per access profile, or, for DHCP only, per address pool. You can configure more than one name server in a routing instance or access profile by repeating the statement for each address.

Because you can configure name server addresses at more than one level, the address returned to the client is determined by the order of preference among the levels. The preference depends on the client type.

- For DHCP subscribers, the preference in descending order is
RADIUS > DHCP address pool > access profile > global
- For non-DHCP subscribers, the preference in descending order is
RADIUS > access profile > global

According to the preference order, a name server address configured in RADIUS is preferred by all subscriber types over all other configuration levels. For all subscriber types, the global name server address is used only when no other name server addresses are configured. When a name server address is configured only in a DHCP address pool, then no address is available to non-DHCP subscribers.

When you configure multiple addresses for a name server, the order in which you configure them determines the preference within that configuration. The preference according to configuration level supersedes this ordering.

There is no restriction on the number of DNS name server addresses that you can configure. For DHCP subscribers, all the addresses are sent in DHCP messages. However, only two addresses—determined by preference order—are sent to PPP subscribers.

All changes in these locally configured DNS name servers affect only new subscribers that subsequently log in. Existing subscribers are not affected by the changes.

Benefits of Local DNS Server Addresses

- Enables configuration of multiple name server addresses per routing instance and per access profile, providing opportunities for subscribers to connect when a given server is unavailable. The multiple

server/multiple level configuration provides a high degree of granularity for managing subscriber access, which is made easier with the capability of specifying a preference order for the servers.

- Supports many subscriber types: Terminated and tunneled PPP subscribers (IPv4 and IPv6), DHCP subscribers (DHCPv4 and DHCPv6), and IP-over-Ethernet (VLAN) subscribers.

SEE ALSO

[Attributes That Can Be Applied to DHCP Clients | 1095](#)

[Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 1100](#)

Configuring DNS Name Server Addresses for Subscriber Management

This topic describes the procedure for configuring DNS name server addresses at the access profile and routing instance levels. For information about configuring addresses in DHCP address pools, see ["Address-Assignment Pools for Subscriber Management" on page 1476](#). For information about configuring addresses on your RADIUS server, refer to your RADIUS software documentation. The order in which the name server configurations at different levels are preferred is described in ["DNS Name Server Address Overview" on page 1503](#).



BEST PRACTICE: In practice, choose either the `domain-name-server` statement or the `domain-name-server-inet` statement for IPv4 addresses. They both have the same effect and there is no need to use both statements. If you do use both statements, addresses configured with `domain-name-server-inet` are preferred over addresses configured with `domain-name-server`.

For example, the following sample configuration specifies two IPv4 domain name servers. The server configured with the `domain-name-server-inet` statement, 192.0.2.23, is preferred over the server configured with the `domain-name-server` statement, 198.51.100.31.

```
[edit access]
user@host# set domain-name-server 198.51.100.31
user@host# set domain-name-server-inet 192.0.2.23
```

To configure DNS name server addresses globally:

- Configure an IPv4 address.

```
[edit access]
user@host# set domain-name-server-inet dns-address
```

Alternatively, you can use a different statement to configure an IPv4 address.

```
[edit access]
user@host# set domain-name-server dns-address
```

- Configure an IPv6 address.

```
[edit access]
user@host# set domain-name-server-inet6 dns-address
```

For example, to configure multiple addresses of each type:

```
[edit access]
user@host# set domain-name-server-inet 198.51.100.31
user@host# set domain-name-server-inet 198.51.100.100
user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:81ca
user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:7334
```

To configure DNS name server addresses in an access profile:

- Configure an IPv4 address.

```
[edit access profile profile-name]
user@host# set domain-name-server-inet dns-address
```

Alternatively, you can use a different statement to configure an IPv4 address.

```
[edit access profile profile-name]
user@host# set domain-name-server dns-address
```

- Configure an IPv6 address.

```
[edit access profile profile-name]
user@host# set domain-name-server-inet6 dns-address
```

For example, to configure multiple addresses of each type:

```
[edit access profile vrf-s-access]
user@host# set domain-name-server-inet 198.51.100.01
user@host# set domain-name-server-inet 198.51.100.100
user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:ac81
user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:71bfd
```

SEE ALSO

[Attributes That Can Be Applied to DHCP Clients | 1095](#)

[Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 1100](#)

Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment

In a DHCPv6 environment, DHCPv6 clients can use a single Solicit message to request multiple addresses (an IA_NA address, an IA_PD address, or both), as well as the DNS server address (DHCPv6 attribute 23). By default, the DHCPv6 local server returns the DNS server address as a global DHCPv6 option.

You can override the default behavior and specify that the DHCPv6 local server returns DNS server addresses as their respective IA_NA and IA_PD suboptions. You can configure the DHCPv6 local server to support the override globally, for a specific group, or for a specific interface.



CAUTION: Some customer premises equipment (CPE) cannot recognize the DNS server address when the address is returned as an IA_NA or IA_PD suboption, which can create interoperability issues.

To configure the DHCPv6 local server to return the DNS server address as an IA_NA or IA_PD suboption.

1. Specify that you want to configure DHCPv6 override options.

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

2. Override the default behavior. DHCPv6 local server now returns DNS server addresses as the respective IA_PD or IA_NA suboption.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set multi-address-embedded-option-response
```

SEE ALSO

[Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 1269](#)

[DHCPv6 Options in a DHCPv6 Multiple Address Environment | 1271](#)

DNS Resolver for IPv6 DNS Overview

In a network that uses Neighbor Discovery Router Advertisement (NDRA) to provide IPv6 addressing, the DNS server address can be provided in Router Advertisements sent to IPv6 hosts. The address is included in a field called Recursive DNS Server (RDNSS). This feature is useful in networks that are not running DHCPv6.

RADIUS can populate the RDNSS field dynamically when an IPv6 subscriber logs in. On the RADIUS server, you can configure a primary and secondary DNS address in the following VSAs, which are stored in the `$junos-ipv6-dns-server` variable:

- Ipv6-Primary-DNS (26-47)
- Ipv6-Secondary-DNS (26-48)

When a subscriber logs in, RADIUS provides the actual DNS server address in the Access-Accept message.

You can also configure a static IPv6 address for DNS servers.

After the subscriber session is established, the DNS address is stored in the session database. When the router sends IPv6 router advertisements, it uses this DNS address in the RDNSS field in the Router Advertisement option.

Configuring a DNS Server Address for IPv6 Hosts

To configure a dynamic DNS server address for IPv6 hosts:

1. Specify that the router receives the DNS server address in the `$junos-ipv6-dns-server-address` variable sent from RADIUS servers in the Access-Accept message when the subscriber logs in.

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-name]
user@host# set dns-server-address $junos-ipv6-dns-server-address
```

2. Specify the time in seconds for which the DNS server address remains valid.

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-name dns-server-address address]
user@host# set lifetime 2400
```

The default value of the lifetime is 1800 seconds.

To configure a static DNS server address for IPv6 hosts:

1. Specify the IPv6 address of the DNS server.

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-name]
user@host# set dns-server-address ipv6-address
```

2. Specify the time in seconds for which the DNS server address remains valid.

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-name dns-server-address address]
user@host# set lifetime 2400
```

The default value of the lifetime is 1800 seconds.

8

PART

M:N Subscriber Redundancy

- [M:N Subscriber Redundancy | 1511](#)
-

CHAPTER 11

M:N Subscriber Redundancy

IN THIS CHAPTER

- [M:N Subscriber Redundancy on BNG | 1511](#)
- [M:N Subscriber Service Redundancy on DHCP Server | 1557](#)
- [N+1 Support for BNG M:N Subscriber Service Redundancy | 1561](#)
- [BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | 1565](#)

M:N Subscriber Redundancy on BNG

IN THIS SECTION

- [M:N Subscriber Redundancy on BNG Overview | 1511](#)
- [How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization | 1543](#)
- [How to Configure M:N Subscriber Redundancy with Pseudowires and DHCP Binding Synchronization | 1549](#)
- [Verifying M:N Redundancy and Active Leasequery Topology Discovery Information | 1554](#)

M:N Subscriber Redundancy on BNG Overview

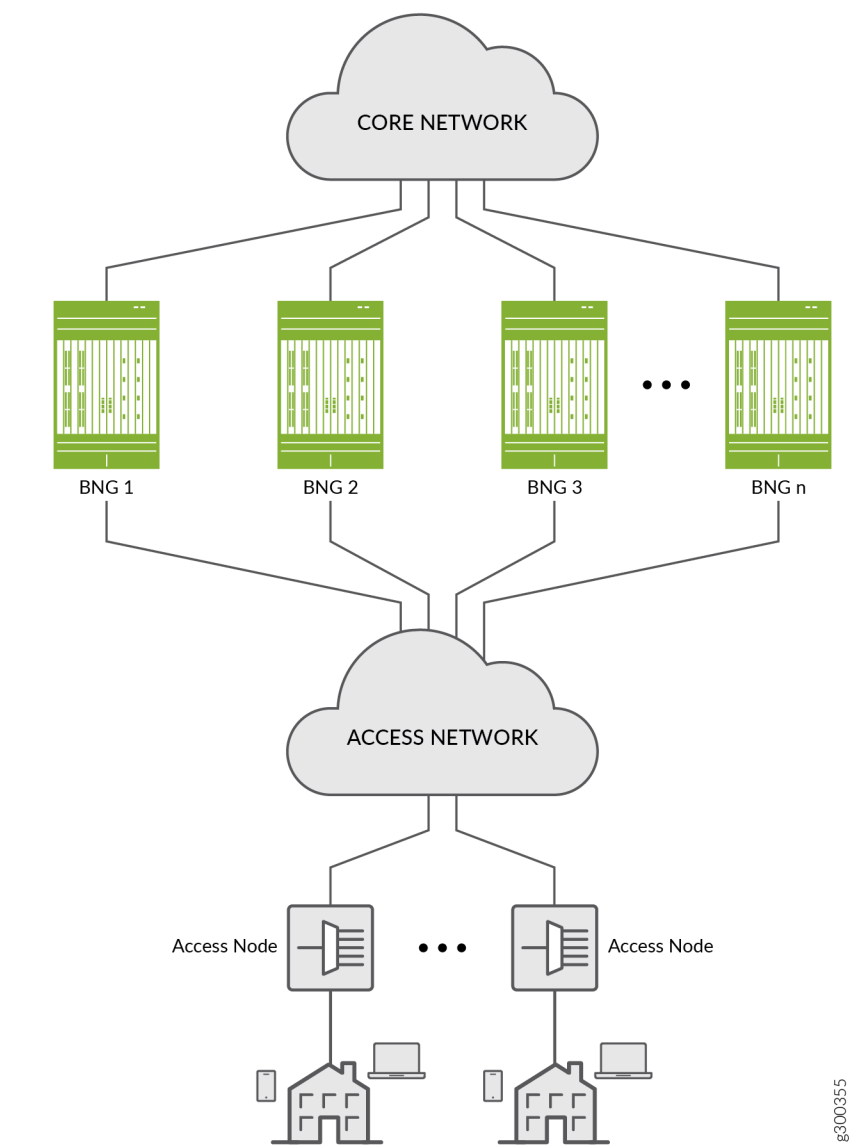
IN THIS SECTION

- [Benefits of M:N Subscriber Redundancy on BNG | 1514](#)
- [Fundamentals of M:N Redundancy | 1515](#)
- [Subscriber Sessions and Hot Standby Mode | 1518](#)
- [M:N Redundancy Using Virtual Router Redundancy Protocol \(VRRP\) | 1519](#)

- VRRP Failover and Reversion Timing | 1522
- M:N Redundancy Using Pseudowire Redundancy | 1523
- DHCP Active Leasequery Topology Discovery and M:N Subscriber Redundancy | 1525
- Example Topology Discovery with VRRP Redundancy | 1527
- Example Topology Discovery with Pseudowire Redundancy | 1530
- Static Subscribers and M:N Redundancy | 1534
- Convergence and M:N Subscriber Redundancy | 1540

Starting in Junos OS Release 19.2R1, you can configure M:N subscriber redundancy as a mechanism for improving network resiliency by protecting subscribers from a variety of software and hardware failures. This protection is available in a typical network topology, like the one shown in [Figure 28 on page 1513](#).

Figure 28: Sample Topology for M:N Subscriber Group Redundancy



A failure in any of the locations listed in [Table 62 on page 1513](#) can trigger a primary BNG to fail over to a backup BNG.

Table 62: Types of Failures Mitigated by M:N Subscriber Group Redundancy

Access line card	Core-facing link
Access link	Partial access network

You can use M:N redundancy to protect the following subscriber types:

- Dynamic DHCPv4 and DHCPv6 subscribers on static 1:1 VLANs over IPoE; VRRP redundancy
- VLAN-based static subscribers; VRRP redundancy
- IP demux-based static subscribers; VRRP redundancy
- DHCPv4 and DHCPv6 subscribers on dynamic or static VLANs over IP/MPLS; pseudowire redundancy (This support is added in Junos OS Release 20.1R1.)

Benefits of M:N Subscriber Redundancy on BNG

- Provides a lightweight, application-layer subscriber redundancy. You can use it to back up multiple different subscriber groups on multiple different BNG chassis. Each subscriber group has one backups in hot-standby mode.
- Multiple BNGs act as both the active BNG for one or more subscriber redundancy groups and as the backup BNG for other subscriber redundancy groups at the same time.
- M:N redundancy is complementary to MX Series Virtual Chassis redundancy. M:N redundancy is appropriate for distributed environments. MX Series Virtual Chassis, requires a dedicated chassis for redundancy. It provides 1:1 redundancy and is most often used in centralized deployments.
- M:N redundancy with DHCP active leasequery topology discovery protects subscribers from several different hardware and software single points of failure. These include failures in access (subscriber-facing) or core-facing links and in an access interface module or the chassis. It also protects against partial access network and partial core network failures.
- You can enable or disable M:N redundancy for subscribers that are active. If you remove the redundancy configuration, subscribers that had the configuration remain intact on both the primary and backup BNGs.
- You can deploy M:N redundancy with a single core-facing interface. This means that multiple subscriber redundancy groups can share a common core connectivity.
- M:N redundancy subscribers can coexist with nonredundancy subscribers. This means that you do not have to have BNGs that are dedicated to subscriber redundancy.
- You can configure M:N redundancy subscribers at run time, even after the subscribers are UP. This is useful for software upgrades, because you can migrate subscribers to backup BNGs and then upgrade the software.

Fundamentals of M:N Redundancy



NOTE: For simplicity, most of the explanation of M:N redundancy in this documentation reflects the use of DHCP subscribers on static VLANs.

The basis of M:N redundancy is that multiple (M) subscriber groups on a given BNG chassis can be backed up on multiple (N) different chassis destinations. We refer to these groups as subscriber redundancy groups.


A subscriber group consists of all subscribers that meet the following criteria:

- (Static VLANs) The subscribers belong to a particular static VLAN and use the same logical access interface, such as ge-1/0/10/1. An access device, such as a switch, DSLAM or OLT, aggregates the subscribers into the common VLAN.
- (Dynamic VLANs) The subscribers belong to the same dynamic VLAN and use the same physical access interface, such as ge-1/0/0.
- (Static IP demux) The subscribers all have a source IP address that matches the configured subnet.

When you configure redundancy for a subscriber group, it becomes a subscriber redundancy group. A given subscriber redundancy group uses only one BNG at a time. We call this BNG the primary. For each subscriber redundancy group, only one of the other BNGs acts as a backup in hot-standby mode. When one of the errors listed in [Table 62 on page 1513](#) occurs for the primary BNG, it fails over to the appropriate backup BNG for the affected redundancy group. This backup BNG is now the new primary BNG for that group. All active subscriber sessions for that subscriber redundancy group are maintained across the failover to the backup BNG.

[Figure 29 on page 1516](#) is a conceptual diagram that illustrates the M:N primary/backup relationships. It shows five BNGs in an M:N primary/backup topology where each BNG has a relationship to every other BNG. If BNG 1 is the primary, you can configure BNG 2, 3, 4, and 5 as the backup BNG for different subscriber redundancy groups. If BNG 2 is the primary, you can configure BNG 1, 3, 4, and 5 as the backup BNG, and so on.

Figure 29: Sample Topology for M:N Subscriber Group Redundancy




For M:N redundancy, it is important to understand that you can configure:

- Only one backup BNG for each subscriber redundancy group.
- A BNG to be the backup router for more than one redundancy group.

This means that a given BNG can be simultaneously both the primary router for many redundancy groups and the backup router for many different redundancy groups. When a primary BNG fails, it fails over to the backup router that you configure for each of its redundancy groups. The subscriber sessions for all redundancy groups on the primary BNG are maintained on all the backup BNGs that become new primaries for the groups.

[Figure 30 on page 1517](#) shows a simple configuration of subscriber groups and subscriber redundancy groups on three DHCP relay agents that are hosted on three BNGs. The BNGs might be directly connected to each other or connected over the access or core networks.

Figure 30: Subscriber Redundancy Groups on Multiple BNGs




- Relay agent RA1 is configured for subscriber redundancy groups, SRG 1 and SRG 2, and subscriber group SG A.
- Relay agent RA2 is configured for SRG 2 and SRG 3.
- Relay agent RA3 is configured for SRG 1, SRG 3, and SG B.

Another way of looking at this is that:

- SRG 1 can be active or backed up on RA1 and RA3.
- SRG 2 can be active or backed up on RA1 and RA2.
- SRG 3 can be active or backed up on RA2 and RA3.
- SG A and SG B are not backed up.


Now consider [Figure 31 on page 1518](#), which shows the same topology, but indicates which BNG is primary and which is backup for each redundancy group. The BNG hosting RA 1 is the primary BNG for SRG 1 and SRG 2.

Figure 31: Primary and Backup BNGs for Subscriber Redundancy Groups Before Failover



If this BNG fails, then it fails over to a different backup BNG for SRG 1 and SRG 2, as shown in [Figure 32 on page 1518](#).

Figure 32: Primary and Backup BNGs for Subscriber Redundancy Groups After Failover



- For SRG 1, it fails over to the BNG hosting RA 3. The RA 3 BNG becomes the new primary for SRG 1.
- For SRG 2, it fails over to the BNG hosting RA 2. The RA 2 BNG becomes the new primary for SRG 2.

The failure has no effect on SRG 3.

Subscriber Sessions and Hot Standby Mode

Each backup BNG is in hot-standby mode for its corresponding primary BNG for each subscriber redundancy group on the backup. This means that the backup BNG is ready to take over from the primary BNG immediately and without disruption when a failover occurs. The following behaviors by the primary and backup BNG enable hot-standby mode to work.

- Subscriber bindings and subscriber state are mirrored synchronously to the backup BNG, as are the primary BNG's ARP and neighbor discovery information. Each subscriber is brought up on the backup BNG and its state is Active. Because the subscribers are active simultaneously on the primary and backup BNG, the backup BNG does not perform any subscriber processing during a failover event.
- Each subscriber session is treated as a continuous session before, during, and after a failover. During initial subscriber login, the primary and backup BNGs each send a RADIUS Accounting-Start message or OCS CCR-I message for the subscriber.

During failover, the failing primary sends an Accounting-Stop or CCR-T message on a best-effort basis. For example, it sends the message if the core-facing link is still up or if the chassis is still running. If the core-facing link is down or the entire chassis is down, then the failing primary can't send an Accounting-Stop or CCR-T message.

When the backup BNG becomes primary, it does not send an Accounting-Start or CCR-I message because the subscribers are active across the failover. Accounting statistics increment from the new primary.

- During initial subscriber login, the BNG adds subscriber routes to its routing table and propagates the routes to the core network. When the primary BNG fails over, it does not delete subscriber routes from its own routing table and it does not withdraw the routes from the core network. After failover, the failed primary does not add or propagate any routes. Alternatively, you can configure the subscriber routes to be advertised to or withdrawn from the core based on BNG primary role so that there is no traffic loss as a result of the failover.



NOTE: State synchronization applies only to subscriber state. Service state is not synchronized. Depending on your services configuration, the BNG might attach services for the subscribers on both active and backup subscribers. Alternatively, the services can reattach after failover on the new active BNG.



NOTE: M:N subscriber redundancy does not synchronize accounting statistics from the primary BNG to the backup BNG. It does make a best-effort attempt to communicate accounting information to an accounting server. When a failover occurs, accounting statistics begin incrementing from the new primary and stop incrementing from the failed primary. Depending on the severity of the failure, failovers can result in loss of accounting information.

M:N Redundancy Using Virtual Router Redundancy Protocol (VRRP)

You can use VRRP to provide M:N redundancy in a network. M:N redundancy uses VRRP to provide a virtual IP address and MAC address shared by two BNGs in a VRRP group (sometimes referred to as a VRRP instance). The VRRP group corresponds to a single virtual router. You configure the VRRP group

on the respective access interface on each BNG. The access interface is the subscriber-facing logical interface that is connected to the access network.

The virtual IP address becomes the default gateway address for the BNGs in the group. Only the BNG acting as the primary sends VRRP advertisements or responds to traffic destined for the virtual router address. The BNG advertises only the virtual gateway address and virtual MAC address to subscriber hosts. Because both routers in the group share the same virtual gateway address, no interaction with the hosts is required and failover from primary to backup occurs within a few seconds.



NOTE: The VRRP solution for M:N redundancy is targeted for an N:1 subscriber access model that uses static underlying logical interfaces.

For detailed information about how VRRP works in general, see [Understanding VRRP](#) and related topic in the *High Availability User Guide*.

You configure different priorities for the two routers in a VRRP group to determine which router the group elects to be the primary:

1. The router with the higher priority for the group is the primary. The larger the number, the higher the priority. For example, between two group members with priorities of 100 and 50, respectively, the router with priority 100 is the primary.
2. When the primary fails, the protocol elects the backup router as the new primary. The new primary assumes ownership of the virtual IP and MAC addresses. Failover has no effect on data traffic.
3. When the original primary comes back online, the protocol determines that it has a higher priority than the current primary (previous backup). The original primary then resumes the primary role with no effect on data traffic.



NOTE: When using VRRP for M:N Subscriber Redundancy, the number of subscriber redundancy groups are limited to the number of supported VRRP sessions on the device. For dual-stack this feature requires separate VRRP sessions for IPv4 and IPv6, therefore the number of subscriber redundancy groups are halved.

[Figure 33 on page 1521](#) shows a sample topology with two BNGs and the configuration for the corresponding interfaces on each router:

- The two logical interfaces are on the same VLAN (1).
- The interface addresses are in the same subnet (203.0.113.1/24 and 203.0.113.2/24).
- The interface addresses are in the same VRRP group (27) and share the same virtual IP address (203.0.113.25).

- The BNG with the higher priority (254) is elected primary; the BNG with the lower priority (200) is the backup.

Figure 33: VRRP Topology and Configuration for Primary and Backup Routers

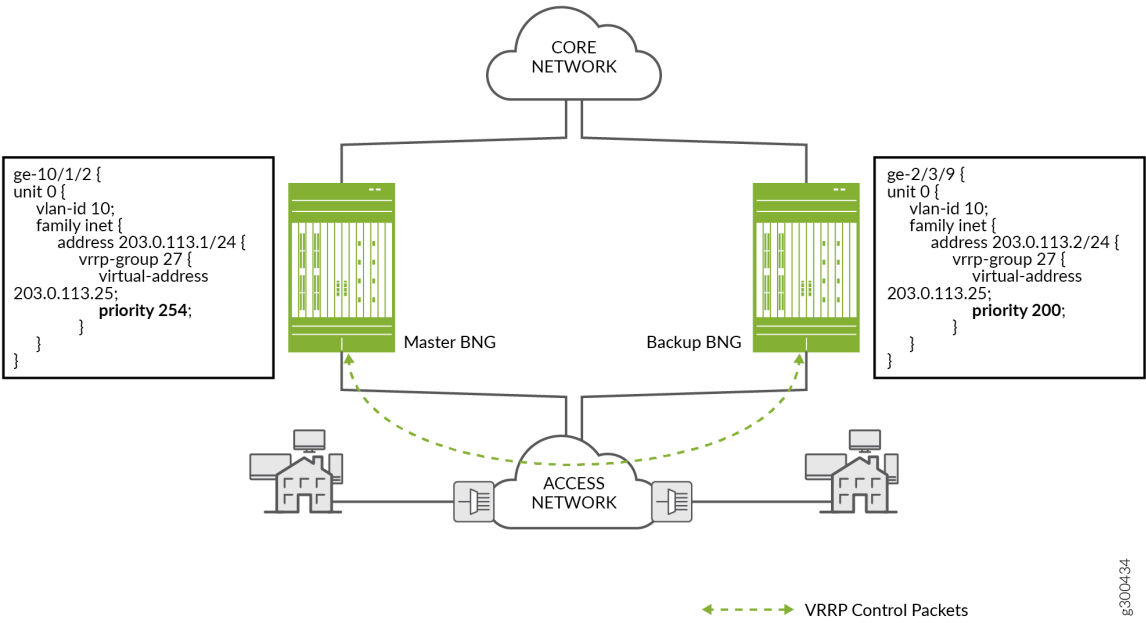
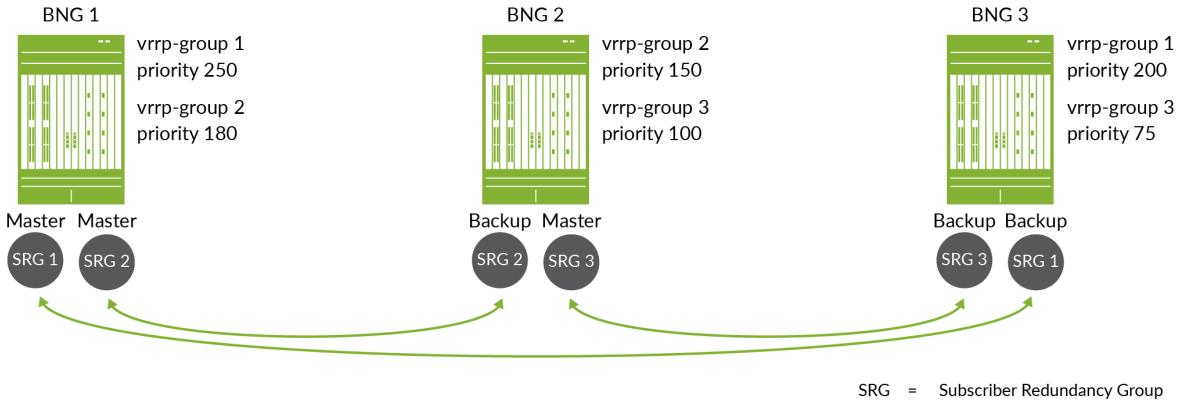


Figure 34 on page 1521 shows how the configured VRRP priority determines which BNG acts as the primary or backup for a subscriber redundancy group.

Figure 34: VRRP Priorities for Three Subscriber Redundancy Groups



The topology includes three subscriber redundancy groups (*M*), SRG 1, SRG 2, and SRG 3 on three BNGs (*N*). Each subscriber redundancy group corresponds to a different VRRP group. The arrows indicate the primary router and backup router for each group

- For SRG 1, BNG 1 has the higher priority, 250. BNG 3 has a lower priority, 200. This means that BNG 1 is the primary for SRG 1 and BNG 3 is the backup, so BNG 1 fails over to BNG 3. When BNG 1 recovers, it is reelected primary for SRG 1, because it has a higher priority than BNG 3.
- For SRG 2, BNG 1 also has the higher priority, 180, and is the primary. BNG 2 has a lower priority, 150, and is the backup.
- For SRG 3, BNG 2 has the higher priority, 100, and is the primary. BNG 3 has a lower priority, 75, and is the backup.

VRRP Failover and Reversion Timing

Using the redundancy configuration shown in [Figure 34 on page 1521](#), suppose BNG 1 fails over to BNG 3 for SRG 1, so that BNG 3 is the new primary for the group. The primary role reverts automatically to BNG 1 when it comes back up. If the connection between the two BNGs is across the access network (as compared to a direct link between the BNGs), the subscriber states might not be synchronized between the two BNGs when the primary role reverts. VRRP state is independent of DHCP active leasequery synchronization.

When the access link on BNG 1 is restored, the DHCP active leasequery restores the connection for subscriber synchronization between the BNGs. DHCP begins to resynchronize the subscriber state and binding information from the current primary (BNG 3) to the recovered original primary (BNG 1).

Accounting statistics can be affected if the primary role reverts to BNG 1 before the resynchronization completes. For example, accounting statistics for subscribers logging in are not added to the database until resynchronization completes. Logout messages for subscribers logging out are not processed until the synch is over and the subscribers are recovered on BNG 1.

You can mitigate these effects by configuring the VRRP hold timer (sometime called the revertive timer) so that resynchronization completes before the original primary resumes the primary role. Use the `hold-time` statement at the `[edit interfaces]` hierarchy level.



BEST PRACTICE: We recommend that you configure VRRP redundancy in non-revertive mode when you are operating at a high scale. For systems not operating at scale, you can either use non-revertive mode or configure the VRRP hold timer (sometime called the revertive timer) with values high enough that resynchronization completes before the original primary resumes the primary role.

M:N Redundancy Using Pseudowire Redundancy

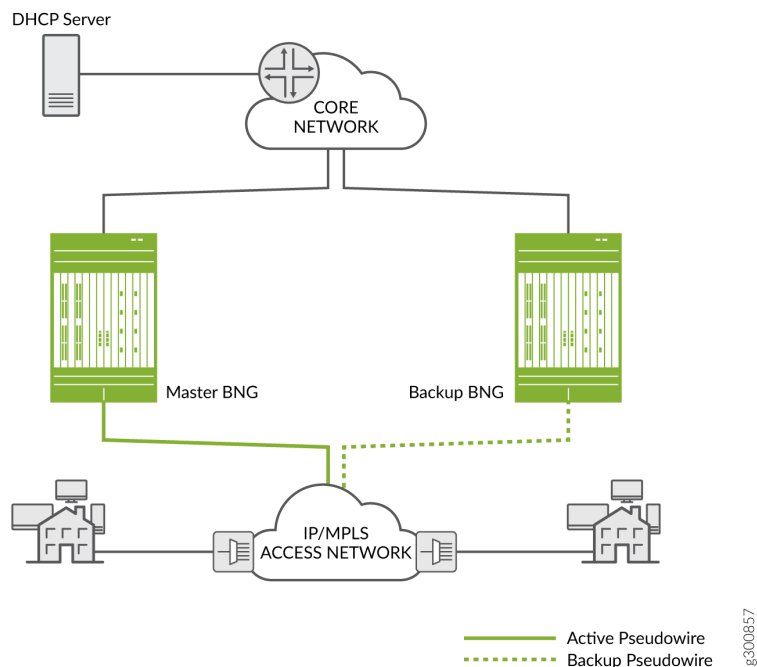
Starting in Junos OS Release 20.1R1, you can use pseudowire redundancy to provide M:N redundancy when the access network consists of Layer 2 (L2) circuits over IP/MPLS. In this type of access network, LDP is the signaling protocol that distributes labels between L2 circuit neighbors. Each L2 circuit is a point-to-point pseudowire tunnel between the access node (or customer edge device) and a BNG. The network can include a heterogeneous mix of L2 or L3 devices.



NOTE: When using Pseudowire Redundancy for M:N Subscriber Redundancy, the number of subscriber redundancy groups are limited to the number of supported pseudowire subscriber interfaces on the device.

Figure 35 on page 1523 shows a simple topology where access nodes aggregate traffic and send it across the network to a DHCP relay agent on the primary BNG. The pseudowire redundancy configuration specifies an active pseudowire (to the primary BNG) and a backup pseudowire (to the backup BNG).

Figure 35: Layer 2 Circuit Topology for Primary and Backup Routers



For L2 circuits, you configure the pseudowires as the underlying (access-facing) interfaces on the BNGs. You then configure the interfaces with L2 connections such as Ethernet, dynamic auto-sensed VLANs, or static VLANs. The DHCP client-facing, pseudowire interfaces are bundled and added to a L2 circuit (the pseudowire tunnel). Typically the bundle includes a set of dynamic VLAN interfaces. However, the

bundle can include any combination of single VLAN logical interfaces, lists of VLAN interfaces, and physical interfaces.

An L2 circuit runs between two L2 neighbors; in this case between an access node and a BNG. Each neighbor serves as an end-point destination for an MPLS label-switched path (LSP). You construct the circuit by configuring it on an interface on each neighbor:

- On the BNG, you specify the access node as a neighbor and a local pseudowire interface on the BNG that ends the L2 circuit.
- On the access node, you specify the BNG as a neighbor and a local interface facing clients on the node that is the other end of the L2 circuit.
- On both the BNG and access node, you configure a unique virtual circuit identifier (VCI) that distinguishes that L2 circuit from among all the other L2 circuits ending on the device.

That L2 circuit is now the primary pseudowire to the BNG. To establish redundancy, you configure the backup pseudowire on the access node. On the same local interface, you specify another BNG as the backup neighbor and specify that the backup pseudowire is in hot-standby mode.

The hot-standby mode ensures that the backup neighbor is fully ready to take over as primary if the current primary circuit fails. An LSP to the backup neighbor is already established by LDP.

The state of the pseudowire interface is UP on the primary BNG. The state of the pseudowire interface is remote standby (RS) on the backup BNG. (You can use the `show l2circuit connections brief` command to view the circuit state.) You must configure your route policies so that subnet routes for this redundancy group are advertised only on the primary BNG. This ensures that only the primary receives downstream traffic.

LDP has a keepalive mechanism to detect failures. A failure results in the L2 circuit failing over from the primary pseudowire and primary BNG to the backup pseudowire and the backup BNG. When it detects a failure, LDP switches the circuit over from the primary LSP (on the primary pseudowire) to the backup LSP (on the backup pseudowire). The backup BNG assumes the primary role and its state transitions to Up.

When the old primary is up again, the same considerations regarding synchronization apply for pseudowire redundancy as they do when VRRP is the redundancy method.



BEST PRACTICE: We recommend that you configure pseudowire redundancy in non-revertive mode when you are operating at a high scale. For systems not operating at scale, you can either use non-revertive mode or configure the `revert-time` interval on the access node interface with values high enough that resynchronization completes before the original primary resumes the primary role.

DHCP Active Leasequery Topology Discovery and M:N Subscriber Redundancy

For DHCP subscribers, DHCP active leasequery and topology discovery enable subscriber state and binding information to be synchronized between peer DHCP relay agents for all subscriber redundancy groups on the peers. This enables leases and data traffic to continue without interruption both when the primary BNG fails over to the backup and when it resumes the primary role.

Although you configure interface-level primary/backup redundancy for pairs of BNGs, it also corresponds in a way to the DHCP relay agents hosted on the primary and backup BNGs. You can think of the DHCP relay agent on the primary BNG as being the primary relay agent for a subscriber redundancy group. Similarly, you can think of the DHCP relay agent on the backup BNG for a group as being the backup relay agent for the group.

Each relay agent that you configure with topology discovery exchanges messages with its configured active leasequery peers to determine the name of access interfaces on its relay agent peers that correspond to and connect with its own local access interfaces. The access interfaces are the interfaces used by the subscriber redundancy groups.

1. When a relay agent sends a topology discovery query message to a peer, that message includes DHCP options that specify the access interface name (Agent Circuit ID), the subnet/mask for the interface, and the VLAN ID for the redundancy group. DHCP also generates a random transaction ID for the exchange that is conveyed in the packet header. The transaction ID is unique for that access interface.
2. The receiving peer relay agent uses the subnet/mask and the VLAN ID to determine whether it has a local access interface for those values. If it does, the peer sends a topology discovery reply over that interface to the querying relay agent's access interface. The reply message includes the subnet/mask, VLAN ID, and the transaction ID that it received in the query.
3. The querying relay agent verifies that the transaction ID in the reply matches the access interface where it received the reply. The transaction ID in the reply must correspond to the one that it sent to the peer for that access interface. If the transaction ID matches, the relay agent can then add an entry to its translation table to associate the two linked interfaces.
4. The querying agent repeats this process for each of its local access interfaces.

[Figure 36 on page 1526](#) shows this query and response for two BNGs when you use VRRP redundancy. BNG 1 sends the query for its access interface, ge-10/1/2, to BNG 2 over the TCP connection. BNG 2 responds over the UDP connection from its associated interface, ge-2/3/9.

BNG 2 sends a query for its access interface to BNG 1 over the TCP connection. BNG 1 responds over the UDP connection from its associated interface ge-10/1/2.

Figure 36: Query and Response for Topology Discovery with VRRP Redundancy

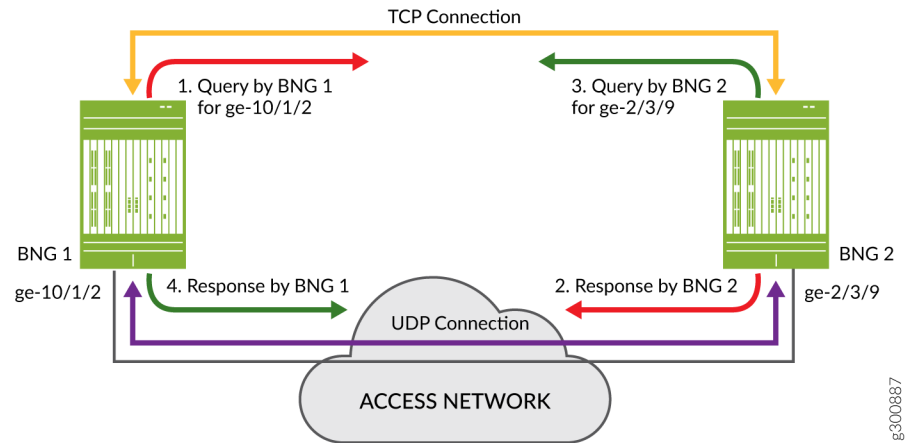
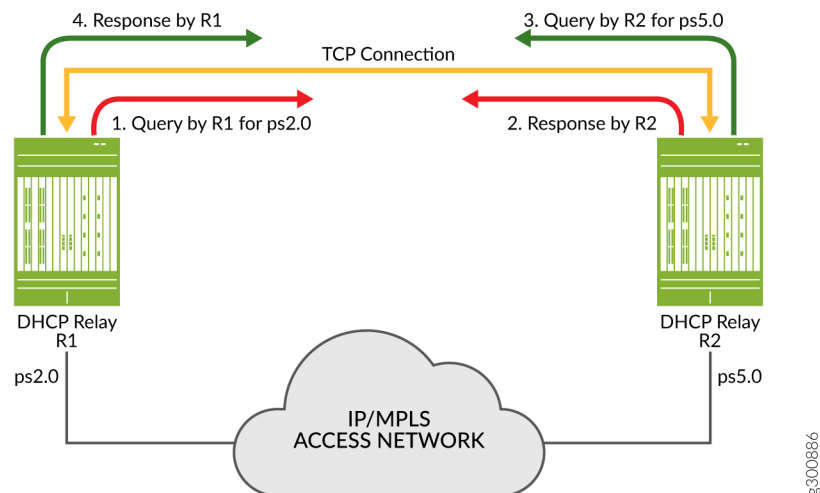


Figure 37 on page 1526 shows a query and response for two DHCP relay agents on BNGs when you use pseudowire redundancy. R1 sends the query for its access interface, ps2.0, to BNG 2 over the TCP connection. R2 responds over the same TCP connection. R2 also sends a query to R1, for its access interface, ps5.0. R1 then responds to this query over the TCP connection. Topology discovery for pseudowire redundancy uses a statically configured, shared common key across BNG pairs as the matching criteria. This is in contrast to VRRP redundancy where matching is performed on subnet/mask and VLAN ID.

Figure 37: Query and Response for Topology Discovery with Pseudowire Redundancy



Each peer agent sends queries to its peers so it can build its own translation table of corresponding local and remote access interfaces. In this way all relay agents that you configure both as peers and for

topology discovery learn the complete set of remote access interfaces for their local interfaces. The translation tables enable the peers to synchronize subscriber information appropriately for each subscriber redundancy group.

After topology discovery is completed, active leasequery performs the subscriber synchronization. Active leasequery performs its queries by giaddr (DHCPv4) or linkaddr (DHCPv6). This query type ensures that DHCP synchronizes only the information for subscribers in a redundancy group for each interface.

You cannot configure this query type; it is a function of configuring topology discovery. When you configure topology discovery, the presence of query-by-relay-id and giaddr in DHCPv4 option 82 or linkaddr in DHCPv6 Option 18 is interpreted to be a query by giaddr or a query by linkaddr, respectively.

The relay agent uses the access interface as the value for its gateway IP address (giaddr or linkaddr) field when it sends packets to the local server on behalf of a client. The local server returns the giaddr/linkaddr when it responds to the relay agent. The relay agent then uses this value to determine where to send the information downstream. The giaddr/linkaddr shows that the packet has been sent for a particular access logical interface, so the relay agent forwards the information to the DHCP client on that interface.

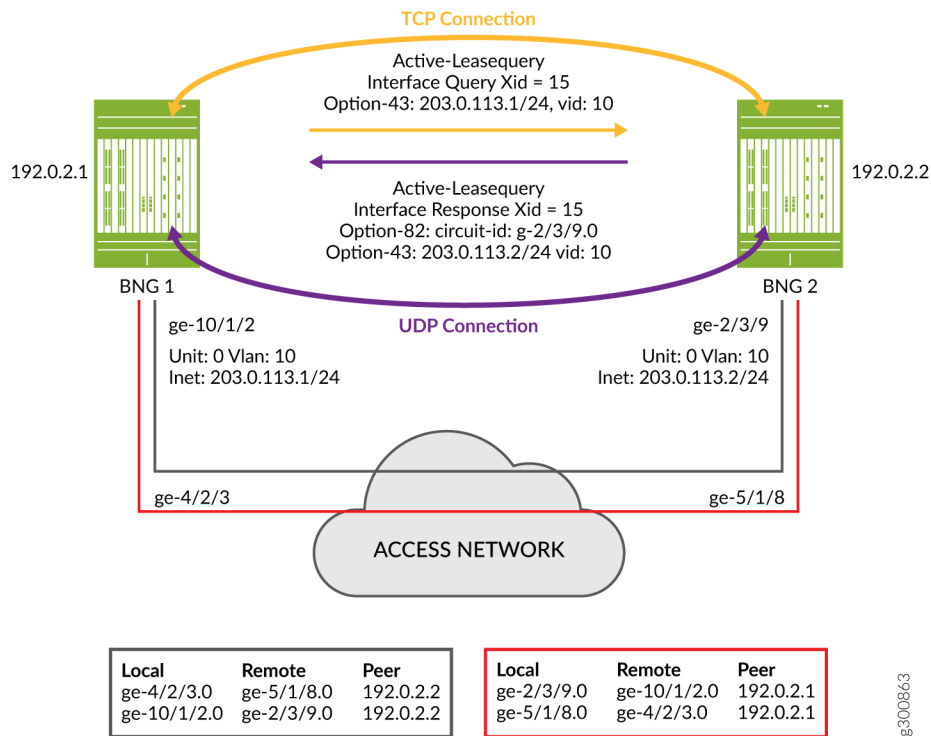
What this means for subscriber redundancy is that by using the giaddr or linkaddr query, active leasequery requests only information for subscribers on that access interface. Consequently, it synchronizes only that subscriber information from the primary relay agent to the backup relay agent. This is a much smaller set of subscribers than if the active leasequery used the query-by-relay-id method, which would return information for all subscribers on the entire chassis.

The result of this process is that each peer agent installs the subscribers for each redundancy group it handles. When the primary BNG/relay agent fails over, the backup already has the necessary subscriber information to maintain the session without interruption.

Example Topology Discovery with VRRP Redundancy

[Figure 38 on page 1528](#) shows a simple topology where active leasequery with topology discovery is configured for the DHCP relay agent peers on two BNGs that are connected over the access network. The configured peer addresses are 192.0.2.1 and 192.0.2.2. We'll use this illustration to understand how topology discovery works when you configure VRRP as the redundancy protocol and how the translation tables are built for each peer relay agent.

Figure 38: Topology Discovery and Translation Tables with VRRP



- After TCP synchronization, peer 192.0.2.1 sends a topology discovery query to peer 192.0.2.2 to determine the matching remote interface for its own local interface, ge-10/1/2.0. Because this is a DHCPv4 topology, the message it sends is a DHCPLEASEQUERY. The query is sent over the TCP connection and includes the following information:
 - The IP subnet address and mask (203.0.113.1/24) of the local access interface, conveyed in DHCPv4 Option 43, suboption 2.
 - The VLAN ID (10) that is configured on the access interface, conveyed in DHCPv4 Option 43, suboption 4.
 - A temporary transaction ID or xid (15), conveyed in the packet header. DHCP generates a random xid for each access interface. The xid is unique across the chassis.

Also included in the query, but not shown in the figure:

 - The client identifier, conveyed in DHCPv4 Option 61.
- Peer 192.0.2.2 receives the query and matches the received subnet address, mask, and VLAN ID to one of its local access interfaces. In this case, the match is to interface ge-2/3/9.0.

3. Peer 192.0.2.2 sends a response back to peer 192.0.2.1 over the UDP connection from its matching access interface, ge-2/3/9.0. The response is a DHCPLEASEACTIVE message and includes the following information:

- The IP subnet address and mask (203.0.113.2/24) of the local access interface, conveyed in DHCPv4 Option 43, suboption 2.
- The VLAN ID (10) that is configured on the access interface, conveyed in DHCPv4 Option 43, suboption 4.
- The name of the matching interface (ge-2/3/9.0), conveyed in Option 82.
- The same temporary transaction ID that it received in the query, conveyed in the IP header.

The following information is also included in the response, but it is not shown in the figure:

- The client identifier, with the same value as that received in the query, in DHCPv4 Option 61.
- The server identifier, in DHCPv4 Option 54.
- The IP destination address in the IP header. This is the subnet address received from peer 192.0.2.1 (203.0.113.1/24).
- The IP source address in the IP header. This is the subnet address (203.0.113.2/24) for this relay agent for the matching interface (ge-2/3/9.0).

4. Peer 192.0.2.1 receives the response over its access interface. It confirms that the transaction ID of the response matches the one it sent in the query. The transaction ID and the vendor-specific suboptions received in the response provide the relay agent with the information it needs to map the two access interfaces in its translation table.

Peer 192.0.2.2 performs the same four steps so that it can update its own translation table. Each of the associated peers initiates topology discovery for all of its local access interfaces. In this way, each peer builds a complete translation table for all of its interfaces.

[Figure 38 on page 1528](#) shows the translation table for each peer that results from the exchange of messages between each pair of peers:

- The relay agent on BNG 1 initiates topology discovery for its three access interfaces.
- The relay agent on BNG 2 initiates topology discovery for its three access interfaces.
- The relay agent on BNG 3 initiates topology discovery for its two access interfaces.



NOTE: Because the transaction ID is generated for only one access interface, topology discovery is successful even when multiple interfaces share the same subnet and VLAN ID.

For example, suppose two interfaces on peer 192.0.2.2 (ge-2/3/9 and ge-11/0/7) match the subnet and VLAN ID that it received in the query.

This relay agent sends a separate response from each of these interfaces to peer 192.0.2.1's interfaces ge-10/1/2.0 and ge-4/2/3.0. The transaction ID does not match interface ge-4/2/3.0 because the querying peer (192.0.2.1) generated the ID for interface ge-10/1/2.0. Consequently, the querying peer updates its translation table only for interface ge-10/1/2.0.

For detailed information about DHCP active leasequery, topology discovery, and how it works with M:N subscriber redundancy, see ["DHCP Active Leasequery" on page 1127](#) and ["Configuring and Using DHCP Active Leasequery" on page 1145](#). The *Topology Discovery Messages* section in ["DHCP Active Leasequery" on page 1127](#) provides descriptions of the information and options carried in the DHCP query and response messages.


Example Topology Discovery with Pseudowire Redundancy

[Figure 39 on page 1531](#) shows a simple topology where active leasequery with topology discovery is configured for the DHCP relay agent peers on two BNGs that are connected over an IP/MPLS access network. The configured peer addresses are 198.51.100.1 and 198.51.100.5. We'll use this illustration to understand how topology discovery works when the access network uses pseudowire tunnels over the IP/MPLS network. Topology discovery for pseudowire redundancy uses a statically configured, shared common key across BNG pairs as the matching criteria. This is in contrast to VRRP redundancy where matching is performed on subnet/mask and VLAN ID. This example also describes how the translation tables are built for each peer relay agent.



NOTE: The topology shows only a TCP connection, because pseudowire M:N redundancy does not use UDP for topology discovery. In contrast, VRRP M:N redundancy uses both TCP and UDP connections.

Figure 39: Topology Discovery and Translation Tables with Pseudowires and Shared Key



1. After TCP synchronization, peer 198.51.100.1 sends a topology discovery query to peer 198.51.100.5 to determine the matching remote interface for its own local interface, ps2.0. Because this is a DHCPv4 topology, the message it sends is a DHCPLEASEQUERY. The query is sent over the TCP connection and includes the following information:

- The shared common key (PseudoWireKey-100.1) configured on the local interface, conveyed in DHCPv4 Option 43, suboption 6.
- A temporary transaction ID or xid (15), conveyed in the packet header. DHCP generates a random xid for each access interface. The xid is unique across the chassis.

Also included in the query, but not shown in the figure:

- The client identifier, conveyed in DHCPv4 Option 61.

2. Peer 198.51.100.5 receives the query and matches the received shared common key to one of its local access interfaces. In this case, the match is to interface ps5.0.
3. Peer 198.51.100.5 sends a response over the TCP connection back to peer 198.51.100.1. The response is a DHCPLEASEACTIVE message and includes the following information:
 - The shared common key (PseudoWireKey-100.1) that it received in the query, conveyed in DHCPv4 Option 43, suboption 6.

- The same temporary transaction ID that it received in the query, conveyed in the IP header.
- The name of the matching interface (ps5.0), conveyed in Option 82.

The following information is also included in the response, but it is not shown in the figure:

- The client identifier, with the same value as that received in the query, in DHCPv4 Option 61.
 - The server identifier, in DHCPv4 Option 54.
4. Peer 198.51.100.1 receives the response over the in-band TCP connection. It confirms that the transaction ID of the response matches the one it sent in the query. The transaction ID and the vendor-specific suboptions received in the response provide the relay agent with the information it needs to map the two access interfaces (local interface ps2.0 and remote interface ps5.0) in its translation table.


Each of the associated peers in a topology initiates topology discovery for each of its local access interfaces. Each peer uses the same four steps described above to build a complete translation table that maps its local interfaces with peer interfaces. In this example topology, that means:

- The DHCP relay agent (R1) on BNG 1 initiates topology discovery for its two access interfaces, ps2.0 and ps3.0.
- The DHCP relay agent (R2) on BNG 2 initiates topology discovery for its two access interfaces, ps4.0 and ps5.0.

You can see the translation table for each peer that results from the exchange of messages between the pair of peers in [Figure 39 on page 1531](#). The same shared common key is configured on both pseudowire interfaces for each pair. For example, ps2.0 and ps5.0 have the key PseudoWireKey-100.1. Interfaces ps3.0 and ps4.0 share a different key (not shown in the figure).

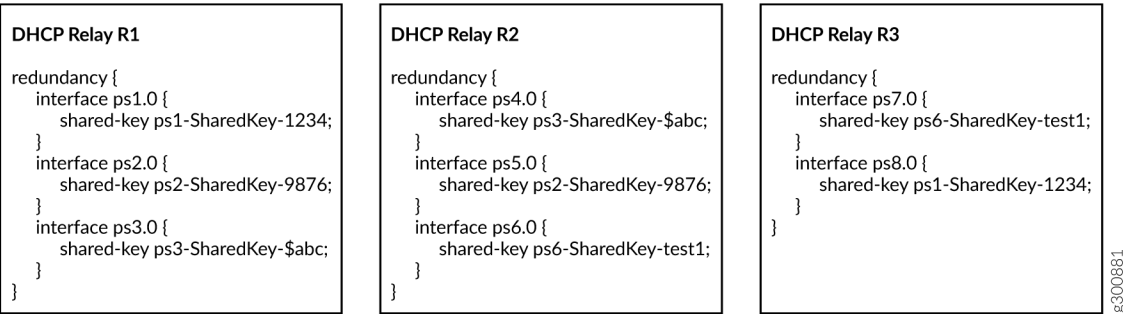
Now consider the slightly more complex topology ,with three peers, shown in [Figure 40 on page 1533](#). Three DHCP relay agents on three BNGs all perform topology discovery for their pseudowire interfaces. The resulting translation tables are shown below each relay agent.

Figure 40: Translation Tables for a Pseudowire Redundancy Topology with Three BNGs



Compare the translation tables and colored pseudowire connection lines in [Figure 40 on page 1533](#) with the shared key configuration snippets for each relay agent in [Figure 41 on page 1533](#).

Figure 41: Sample Shared Key Configuration for Three Peers



You can see that interface ps1.0 on R1 has the same shared key as interface ps8.0 on R3. The translation tables for R1 and R3 show this relationship was discovered by the topology discovery process.

Similarly, interface ps2.0 on R1 and ps5.0 on R2 have the same shared key. Again, topology discovery determined this relationship and each agent updated its translation table accordingly. The other rows in the translation tables were populated in the same way.

For detailed information about DHCP active leasequery, topology discovery, and how it works with M:N subscriber redundancy, see ["DHCP Active Leasequery" on page 1127](#) and ["Configuring and Using DHCP Active Leasequery" on page 1145](#). The *Topology Discovery Messages* section in ["DHCP Active Leasequery" on page 1127](#) provides descriptions of the information and options carried in the DHCPv4 and DHCPv6 query and response messages.

Static Subscribers and M:N Redundancy

M:N subscriber redundancy supports two categories of subscribers:

- Subscribers that use the DHCP client protocol over a static VLAN. This is the most common subscriber type for M:N subscriber redundancy.
- Subscribers on static interfaces that are not running a client protocol. This subscriber type is typical for small to medium enterprises that have their own static IP address and do not use anything like DHCP.

Static subscribers consist of the following types:

- VLAN-based static subscribers—You create subscribers on top of the VLAN logical interface. You configure the VRRP attributes on the VLAN logical interface.
- IP demux-based static subscribers—You create subscribers on an IP demux interface over an underlying interface. Traffic for these subscribers includes a source IP address that matches the configured subnet for the subscriber interface. You configure VRRP attributes on the underlying logical interface.

Both of these static subscriber types are managed by the jsscd daemon. They are sometimes referred to as JSSCD static subscribers.

The following sample configuration snippets show you how to create a static subscriber group with two interfaces configured for VRRP on a primary BNG and a backup BNG. One interface is an IP demux interface and the other is a VLAN interface. The configuration shows how VRRP is configured on each interface.

Primary BNG configuration:

1. The following snippet configures the underlying interface for the IP demux logical interface, ge-1/1/9.11. It specifies the VLAN ID as 11. The access interface subnet is set to 203.0.113.1/24. The VRRP configuration on this subnet sets the group (the subscriber redundancy group) to 11 and specifies the address for the virtual router. The virtual router consists of the primary and backup

BNGs for this subscriber redundancy group. The VRRP priority is 230. When the primary fails over to the backup, assumption of the primary role by the backup is delayed by 30 seconds.

```
[edit]
interfaces {
  ge-1/1/9 {
    unit 11 {
      demux-source inet;
      vlan-id 11;
      family inet {
        address 203.0.113.1/24 {
          vrrp-group 11 {
            virtual-address 203.0.113.25;
            priority 230;
            preempt {
              hold-time 30;
            }
          }
        }
      }
    }
  }
}
```

2. The following snippet configures the VLAN logical interface, ge-1/1/9.20. It specifies the VLAN ID as 20. The access interface subnet is set to 192.0.2.1/24. The VRRP configuration on this subnet sets the group (the subscriber redundancy group) to 20 and specifies the address for the virtual router. The virtual router consists of the primary and backup BNGs for this subscriber redundancy group. The VRRP priority is 230. When the primary fails over to the backup, assumption of the primary role by the backup is delayed by 30 seconds.

```
[edit]
interfaces {
  ge-1/1/9 {
    unit 20 {
      vlan-id 20 ;
      family inet {
        address 192.0.2.1/24 {
          vrrp-group 20 {
            virtual-address 192.0.2.25;
            priority 230;
          }
        }
      }
    }
  }
}
```

```

                                preempt {
                                    hold-time 30;
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}

```

3. The following snippet configures the IP demux logical interface, demux0.1, over the underlying interface, ge-1/1/9.11. It also configures the loopback interface and enables the local address for the IP demux interface to be derived from the loopback interface.

```

[edit]
interfaces {
    demux0 {
        unit 1 {
            demux-options {
                underlying-interface ge-1/1/9.11;
            }
            family inet {
                unnumbered-address lo0.0;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.10.32/32;
            }
        }
    }
}

```

4. The following snippet configures a static subscriber group, static-ifl, that includes both the IP demux static subscriber interface (demux0.1) and the VLAN static subscriber interface (ge-1/1/9.20). It associates an access profile with the group, sets the password and a prefix for the username.

```

[edit system services]
static-subscribers {

```

```

group static-ifl {
    access-profile {
        staticauth;
    }
    authentication {
        password "$ABC123$ABC123"; ## SECRET-DATA
        username-include {
            user-prefix test-static;
        }
    }
    interface ge-1/1/9.20;
    interface demux0.1;
}
}

```

5. The following snippet configures an access profile for the static subscribers group.

```

[edit access]
profile staticauth {
    authentication-order none;
}

```

Backup BNG configuration:



NOTE: In this example, some configuration details are different and others must be the same.

- The access interfaces are different. Alternatively, you can configure the access interfaces to be the same on the primary and backup.
- The VRRP priority is set to 200 for both interfaces. That value makes this the backup BNG, because it is lower than the priority on the other BNG (230).
- The interface addresses are different. The virtual address is the same for both, as it must be, so that both BNGs are in the same virtual router.
- The access interfaces are on the same subnet.

1. The following snippet configures the underlying interface for the IP demux logical interface, ge-3/0/1.11. It specifies the VLAN ID as 11. The access interface subnet is set to 203.0.113.2/24. The VRRP configuration on this subnet sets the group (the subscriber redundancy group) to 11 and specifies the address for the virtual router. The virtual router consists of the primary and backup

BNGs for this subscriber redundancy group. The VRRP priority is 200. When the primary fails over to the backup, assumption of primary role by the backup is delayed by 30 seconds.

```
[edit]
interfaces {
  ge-3/0/1 {
    unit 11 {
      demux-source inet;
      vlan-id 11;
      family inet {
        address 203.0.113.2/24 {
          vrrp-group 11 {
            virtual-address 203.0.113.25;
            priority 200;
            preempt {
              hold-time 30;
            }
          }
        }
      }
    }
  }
}
```

2. The following snippet configures the VLAN logical interface, ge-3/0/1.20. It specifies the VLAN ID as 20. The access interface subnet is set to 192.0.2.2/24. The VRRP configuration on this subnet sets the group (the subscriber redundancy group) to 20 and specifies the address for the virtual router. The virtual router consists of the primary and backup BNGs for this subscriber redundancy group. The VRRP priority is 200. When the primary fails over to the backup, assumption of the primary role by the backup is delayed by 30 seconds.

```
[edit]
interfaces {
  ge-3/0/1 {
    unit 20 {
      vlan-id 20 ;
      family inet {
        address 192.0.2.2/24 {
          vrrp-group 20 {
            virtual-address 192.0.2.25;
            priority 200;
          }
        }
      }
    }
  }
}
```

```

                                preempt {
                                    hold-time 30;
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}

```

3. The following snippet configures the IP demux logical interface, demux0.1, over the underlying interface, ge-3/0/1.11. It also configures the loopback interface and enables the local address for the IP demux interface to be derived from the loopback interface.

```

[edit]
interfaces {
    demux0 {
        unit 1 {
            demux-options {
                underlying-interface ge-3/0/1.11;
            }
            family inet {
                unnumbered-address lo0.0;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.10.32/32;
            }
        }
    }
}

```

4. The following snippet configures a static subscriber group, static-ifl, that includes both the IP demux static subscriber interface (demux0.1) and the VLAN static subscriber interface (ge-3/0/1.20). It associates an access profile with the group, sets the password and a prefix for the username.

```

[edit system services]
static-subscribers {

```

```

group static-ifl {
    access-profile {
        staticauth;
    }
    authentication {
        password "$ABC123"; ## SECRET-DATA
        username-include {
            user-prefix test-static;
        }
    }
    interface ge-3/0/1.20;
    interface demux0.1;
}
}

```

5. The following snippet configures an access profile for the static subscribers group.

```

[edit access]
profile staticauth {
    authentication-order none;
}

```

Convergence and M:N Subscriber Redundancy

Convergence is the process where routers in a network update their individual routing tables when routes on any router are added, removed, or no longer reachable because of a link failure. The routing protocols on the routers advertise the route changes throughout the network. As each router receives the updates, it recalculates the routes and then builds new routing tables based on the results.

A network is *converged* when all the routing tables agree on the overall network topology. For example, this means that the routers have a common understanding about which links are up or down, and so on. How long it takes the routers to reach a state of convergence is called the *convergence time*. The length of the convergence time depends on various factors, such as the size and complexity of the network and the performance of the routing protocols.

M:N subscriber redundancy supports both access-side (upstream) and core-side (downstream) route convergence. Because each subscriber is active simultaneously on the primary and the backup BNGs, traffic convergence can be very quick. However, route convergence is best effort and depends on the degree of failover; that is, whether a partial or complete chassis failure occurs.

It is up to you to determine how to manage upstream and downstream traffic convergence for your network after a failover from primary to backup BNG.

Upstream Traffic Convergence (VRRP Redundancy)

You can improve upstream traffic convergence by using gratuitous ARP to reduce the time it takes for the access network to begin sending traffic to the new primary BNG after the original primary BNG fails.

1. On the primary BNG, the access interface or interface module goes down.
2. VRRP elects the backup BNG as the new primary.
3. The new primary broadcasts gratuitous ARP messages to the access network. It sends the messages from its access interface corresponding to the former primary's access interface. The ARP message contains the VRRP virtual IP address and virtual MAC address that define the virtual router that includes the two BNGs.
4. The switch or other device on the access network relearns the gateway IP address (the virtual address). When it sends traffic to that address, the new primary BNG receives it on the access interface.

Upstream Traffic Convergence (Pseudowire Redundancy)

When you configure the primary and backup pseudowires in hot-standby mode on the access node, LDP automatically establishes LSPs to the primary and backup BNGs. The LDP signaling protocol includes a keepalive mechanism to detect failures in the path. In this case, upstream convergence is achieved by a pseudowire Layer 2 tunnel switch from the primary BNG to the backup BNG.

You can configure LDP keep-alive timers for faster detection of failures. Alternatively, you can run the BFD protocol for faster failover. Any of the following methods can cause a switch from the primary pseudowire to the backup pseudowire:

- Use the request `l2circuit-switchover` command to manually trigger a switch from the primary pseudowire to the backup pseudowire.
- You can configure Bidirectional Forwarding Detection (BFD) for the LDP LSPs. BFD liveness detection can detect two different kinds of failures:
 - A link failure in the LSP path between the access node and the primary BNG. In this case the BNG is still up.
 - A neighbor down failure when the primary BNG goes down.

For both types, you control the speed of the detection and switchover by the configuration of the `bfd-liveness-detection` statement at the `[edit protocols ldp oam]` hierarchy level.

Downstream Traffic Convergence

The time required for downstream traffic convergence is affected by several factors, including the following:

- Advertising individual subscriber routes increases the number of route recalculations that the core network routers must perform.
- Detecting when an access interface goes down and then sending the appropriate route change notification to the core can sometimes be difficult or take a long time.
- Routing protocols at the core might not learn immediately when either a core-facing link or the entire chassis fails. Routing protocols typically rely on some type of timeout to detect the loss, so there is always a delay waiting for the timeout to expire.

We recommend the following guidelines:

- Ensure that subscriber routes are aggregated for advertisement to the core whenever possible. Aggregation might be achieved by using address pools or policy-based route advertisement as described below. Reducing the number of routes to be recalculated on the core routers reduces convergence time, especially as the scale of subscribers increases.
- Configure the routes to be advertised from both BNGs with different preferences. Use fast rerouting techniques at the core.
- Avoid load balancing downstream traffic between the primary and backup BNGs.

Two methods you might consider are policy-based route advertisement and dedicated BNG links.

- Policy-based route advertisement (VRRP and pseudowire redundancy)—This technique can reduce downstream traffic convergence time because only aggregated routes are updated in the core network, rather than numerous individual subscriber routes. For this method, you configure BGP, OSPF, or any other routing protocol to advertise aggregated routes toward the core only when a BNG becomes the primary.

For VRRP redundancy, you configure the BGP policies to track the VRRP virtual IP address. BGP aggregates the subscriber routes based on the subscriber redundancy group corresponding to a VRRP group. BGP advertises the aggregated routes to the core when the VRRP primary role is assumed by the BNG.

For pseudowire redundancy, you configure the BGP policies to track the pseudowire interface status (Up or Down). BGP aggregates routes for the subscriber redundancy group. BGP advertises the aggregated routes to the core when the state changes to Up, meaning that the backup BNG is now the primary.

In either case, if the primary BNG fails over to the backup, BGP on the failed primary withdraws the aggregated subscriber routes for the core. When the backup BNG becomes the new primary, it in turn advertises aggregated subscriber groups to the core.

- **BNG dedicated links (VRRP redundancy only)**—You can reduce the time it takes to detect a failure on the primary BNG by connecting the BNGs with a dedicated link. You configure VRRP on the access interface to track the state of the dedicated link interface. You also configure VRRP on the dedicated link interface to track the state of the access interface.

A failure on the access interface on the primary causes the VRRP primary role to change on the dedicated link. That change in turn causes the primary role to change immediately on the access interface on the backup BNG. This method is faster than waiting for the VRRP hello timer to expire.

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization

IN THIS SECTION

- [Configure Subscriber Group Redundancy | 1544](#)
- [Configure VRRP to Support M:N Redundancy | 1546](#)
- [Configure Active Leasequery with Topology Discovery | 1547](#)

M:N subscriber redundancy with VRRP and DHCP binding synchronization requires you to configure all of the following:

- Redundant subscriber groups to specify the subscribers that are part of the primary/backup operation.
- VRRP on all redundant routers in the topology. VRRP is the protocol that provides the underlying redundancy capability for the subscriber groups and DHCP relay agents.
- DHCP active leasequery with topology discovery for all peer DHCP relay agents in the topology. Active leasequery is responsible for synchronizing the subscriber state and binding information among the peer relay agents. Topology discovery enables the peer relay agents to determine the remote access interfaces for their subscriber redundancy groups so that they can build translation tables of local and remote interfaces to support the M:N primary/backup redundancy scheme.



NOTE: This topic describes only the basic configurations necessary for M:N subscriber redundancy on the BNGs that host the peer DHCP relay agents. It does not describe

every aspect of the following: global subscriber management, the VRRP configuration that you might use in your network, DHCP relay agents, or DHCP leasequery. For more information about these subjects, see the following:

- *Junos OS Enhanced Subscriber Management* and *Configuring Junos OS Enhanced Subscriber Management*
- ["DHCP Leasequery Methods" on page 1116](#)
- [Understanding VRRP](#) and [Configuring VRRP](#)



NOTE: M:N subscriber redundancy requires that the primary and backup BNGs support the same protocol versions for DHCP and VRRP. If the protocol support is different between the BNGs, you might see undesirable side-effects.



NOTE: Dual-stack redundancy subscribers have the following requirements:

- DHCP configuration—You must configure active leasequery with topology discovery for both DHCPv4 and DHCPv6.
- VRRP configuration—You must configure both address families on the access interface, because dual-stack subscribers require two sessions, one each for IPv4 and IPv6. You must also configure the same VRRP primary role priority for the IPv4 and IPv6 sessions for a given redundancy group because they share the same logical interface.

Configure Subscriber Group Redundancy

To configure subscriber group redundancy on a BNG:

1. Access the redundancy stanza.

```
[edit system services subscriber-management]
user@host# edit redundancy
```

2. (Optional) Specify VRRP as the redundancy method.



NOTE: This value is set by default.

```
[edit system services subscriber-management redundancy]
user@host# set protocol vrrp
```

3. Specify the names of the access interfaces used by subscribers that you want in redundancy groups. You must specify all such interfaces that are on the chassis, regardless of how you later organize them into redundancy groups.



NOTE: Only Gigabit Ethernet (ge) and 10-Gigabit Ethernet (xe) interfaces are supported.

```
[edit system services subscriber-management redundancy]
user@host# set interface name1
user@host# set interface name2
...
```

4. Configure IPv4, IPv6, or both IPv4 and IPv6 virtual addresses. You must configure both families for dual-stack subscribers. VRRP uses this virtual address to create a virtual router for the BNGs that support a particular subscriber redundancy group. This means that you must configure the same virtual addresses on each of those BNGs.

```
[edit system services subscriber-management redundancy]
user@host# set interface name virtual-inet-address virtual-inet-address
user@host# set interface name virtual-inet6-address virtual-inet6-address
```

5. (Optional) Suppress advertisement of subscriber access routes or framed routes at the backup BNG towards the core or install the routes in the forwarding table. The routes are added to the routing table when the primary fails over to the backup BNG. This option applies to all subscribers that are covered by redundancy on the chassis and that log in after you configure the option. Existing subscribers are not affected.



BEST PRACTICE: We recommend that you always configure `no-advertise-routes-on-backup` when you use the non-aggregated mode of address allocation. This address-allocation mode increases downstream traffic convergence when a primary BNG fails over to a backup. The `no-advertise-routes-on-backup` option reduces the number of routes advertised and the associated potential issues.

However, we recommend that you use the aggregated mode of address allocation instead, whenever possible. This address-allocation mode enables the fastest downstream traffic convergence if a primary BNG fails over.

```
[edit system services subscriber-management redundancy]
user@host# set no-advertise-routes-on-backup
```

Configure VRRP to Support M:N Redundancy

To configure VRRP to support M:N redundancy for a subscriber redundancy group on a BNG:

1. Configure the logical access interface for the subscriber redundancy group.

```
[edit]
user@host# edit interfaces interface-name unit logical-unit-number
```

2. Specify the VLAN ID common to all members of the subscriber redundancy group (VRRP group).

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set vlan-id vlan-id
```

3. Configure the address family for the access interface.



NOTE: This sample procedure shows only the IPv4 address family, but you might configure the IPv6 address family, or both IPv4 and IPv6. Dual-stack subscribers require two sessions, one each for IPv4 and IPv6, so you must configure both address families on the interface.

```
[edit interfaces interface-name unit logical-unit-number]
user@host# edit family inet
```

4. Specify the subnet (the subscriber-facing address/mask) for the local access interface for the subscriber redundancy group.

```
[edit interfaces interface-name unit logical-unit-number family inet]
user@host# set address address
```

5. Specify the VRRP group identifier. The VRRP group corresponds to a subscriber redundancy group.

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
user@host# set vrrp-group group-id
```

6. Configure the virtual IP address that is used as the default gateway for all BNGs in the same VRRP (subscriber redundancy) group.

This is the same address that you configure with the `virtual-inet-address` or `virtual-inet6-address` options at the `[edit system services subscriber-management redundancy interface]` hierarchy level.

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id]
user@host# set virtual-address address
```

7. Configure the router's priority for becoming the primary router for the redundancy group. A router with a higher number has priority over a router with a lower number.



NOTE: For dual-stack subscribers, you must configure the same priority for the IPv4 and IPv6 sessions for a given redundancy group because they share the same logical interface.

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id]
user@host# set priority number
```

8. (Optional) Configure the hold (revertive) timer to enable subscriber synchronization to complete between BNGs before primary-role reversion completes when the higher-priority primary recovers.

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id]
user@host# set preempt hold-time seconds
```

Configure Active Leasequery with Topology Discovery

Enable active leasequery with topology discovery on the pair of DHCP relay agents that support a given subscriber redundancy group. You must repeat the configuration for each pair of relay agents for different redundancy groups.



NOTE: The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level.



NOTE: For dual-stack subscribers, you must configure active leasequery with topology discovery for both DHCPv4 and DHCPv6.



NOTE: Because active leasequery is an extension of bulk leasequery, you must also configure bulk leasequery for active leasequery to operate. You must configure bulk leasequery before you configure active leasequery. See ["Configuring and Using DHCP Bulk Leasequery" on page 1141](#).

1. Specify that you want to configure active leasequery options for the DHCP relay agent.

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

2. Specify the IP address for a peer with which this relay agent synchronizes information. You must also configure active leasequery on the peer.



NOTE: This is the address used for the TCP connection. It can be a physical interface address or a loopback address per pair of peers.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set peer-address ip-address
```

3. Configure the relay agent to send topology discovery messages to determine the remote access interfaces for subscriber redundancy groups on similarly configured peer relay agents. Discovering the topology enables the relay agents to build translation tables of local and remote interfaces to support an interface-level, primary/backup redundancy scheme.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set topology-discover
```

4. Configure the relay agent to always include Option 82, suboption 1, the Agent Circuit ID. This is the name of the access interface.



NOTE: For DHCPv6, the equivalent statement is *relay-agent-interface-id* to include Option 18.

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

How to Configure M:N Subscriber Redundancy with Pseudowires and DHCP Binding Synchronization

IN THIS SECTION

- [Configure Subscriber Group Redundancy | 1550](#)
- [Configure Active Leasequery with Topology Discovery | 1553](#)

M:N subscriber redundancy with pseudowires and DHCP binding synchronization requires you to configure all of the following:

- Redundant subscriber groups to specify the subscribers that are part of the primary/backup operation.
- DHCP active leasequery with topology discovery for all peer DHCP relay agents in the topology. Active leasequery is responsible for synchronizing the subscriber state and binding information among the peer relay agents. Topology discovery enables the peer relay agents to determine the remote access interfaces for their subscriber redundancy groups so that they can build translation tables of local and remote interfaces to support the M:N primary/backup redundancy scheme.



NOTE: M:N subscriber redundancy with pseudowires functions in an IP/MPLS network where pseudowire tunnels from an access node (such as a switch) make up the L2 circuits to the primary and backup BNGs acting as DHCP relay agents. Those configurations are outside the scope of this documentation.

This topic describes only the basic configurations necessary for M:N subscriber redundancy on the BNGs that host the peer DHCP relay agents. It does not describe every aspect of the following: global subscriber management, DHCP relay agents, or

DHCP leasequery. It does not describe how to configure your IP/MPLS network, the access node that creates the L2 circuits to the DHCP relay agents, or the pseudowire tunnels. For more information about these subjects, see the following:

- *Junos OS Enhanced Subscriber Management and Configuring Junos OS Enhanced Subscriber Management*
- ["DHCP Leasequery Methods" on page 1116](#)
- *MPLS Pseudowire Subscriber Logical Interfaces*
- *Redundant Pseudowires for Layer 2 Circuits and VPLS*
- [Understanding the LDP Signaling Protocol](#)
- [MPLS Applications User Guide](#)



NOTE: M:N subscriber redundancy requires that the primary and backup BNGs support the same protocol versions for DHCP. If the protocol support is different between the BNGs, you might see undesirable side-effects.



NOTE: Dual-stack redundancy subscribers have the following requirement:

- DHCP configuration—You must configure active leasequery with topology discovery for both DHCPv4 and DHCPv6.

Configure Subscriber Group Redundancy

To configure subscriber group redundancy on a BNG:

1. Access the redundancy stanza.

```
[edit system services subscriber-management]
user@host# edit redundancy
```

2. (Optional) Specify pseudowire as the redundancy method.

```
[edit system services subscriber-management redundancy]
user@host# set protocol pseudo-wire
```

3. Specify the names of the pseudowire access interfaces used by subscribers that you want in redundancy groups. You must specify all such interfaces that are on the chassis, regardless of how you later organize them into redundancy groups.



NOTE: Only pseudowire (ps) interfaces are supported.

```
[edit system services subscriber-management redundancy]
user@host# set interface name1
user@host# set interface name2
...
```

4. Configure IPv4, IPv6, or both IPv4 and IPv6 local addresses for the associated pseudowire interface. You must configure both families for dual-stack subscribers. The local IP address must match one of the access-facing GE interface addresses. The local IP address is unique per subscriber redundancy group (identified by the pseudowire `ps.x.0`)

```
[edit system services subscriber-management redundancy]
user@host# set interface name local-inet-address v4-address
user@host# set interface name local-inet6-address v6-address
```

Active leasequery uses this local address as the gateway IP address when it uses the query by giaddr (DHCPv4) or query by linkaddr (DHCPv6) method to query the peer BNG. The relay agent evaluates the giaddr/linkaddr and sends information to the DHCP client that uses the access interface matching the giaddr/linkaddr.

5. Configure the shared common key that identifies the primary and backup pseudowire interfaces on BNG redundancy peers.



NOTE: You must configure a given shared key only on the matching interfaces for a pair of redundancy peers. You must not configure that key on any other peer BNGs.

```
[edit system services subscriber-management redundancy]
user@host# set interface name shared-key string
user@host# set interface name shared-key string
```

6. (Optional) Suppress advertisement of subscriber access routes or framed routes at the backup BNG towards the core or install the routes in the forwarding table. The routes are added to the routing table when the primary fails over to the backup BNG. This option applies to all subscribers that are

covered by redundancy on the chassis and that log in after you configure the option. Existing subscribers are not affected.



BEST PRACTICE: We recommend that you always configure `no-advertise-routes-on-backup` when you use the non-aggregated mode of address allocation. This address-allocation mode increases downstream traffic convergence when a primary BNG fails over to a backup. The `no-advertise-routes-on-backup` option reduces the number of routes advertised and the associated potential issues.

However, we recommend that you use the aggregated mode of address allocation instead, whenever possible. This address-allocation mode enables the fastest downstream traffic convergence if a primary BNG fails over.

```
[edit system services subscriber-management redundancy]
user@host# set no-advertise-routes-on-backup
```

For example, you might configure the following on one BNG:

```
[edit system services subscriber-management redundancy]
user@host# set protocol pseudo-wire
user@host# set interface ps2.0 local-inet-address 10.80.1.2
user@host# set interface ps2.0 local-inet6-address 2001:db8::
user@host# set interface ps2.0 shared-key pskey-2.0-abc-215
user@host# set interface ps3.0 local-inet-address 10.10.0.1
user@host# set interface ps3.0 local-inet6-address 2001:db8:ff:f8::
user@host# set interface ps3.0 shared-key pskey-3.0-def-43
user@host# set no-advertise-routes-on-backup
```

Then configure the following on a peer BNG. Note that `ps5.0` on this BNG shares the same key as `ps2.0` on the other. That signifies that `ps2.0` and `ps5.0` are the associated access interfaces for pseudowire redundancy. Similarly, associated interfaces `ps3.0` and `ps4.0` have the same shared key as each other.

```
[edit system services subscriber-management redundancy]
user@host# set protocol pseudo-wire
user@host# set interface ps4.0 local-inet-address 10.55.3.0
user@host# set interface ps4.0 local-inet6-address 2001:db8:1c:44::
user@host# set interface ps4.0 shared-key pskey-3.0-def-43
user@host# set interface ps5.0 local-inet-address 10.60.20.1
user@host# set interface ps5.0 local-inet6-address 2001:db8:01:10:cd::
```

```
user@host# set interface ps5.0 shared-key pskey-2.0-abc-215
user@host# set no-advertise-routes-on-backup
```

Configure Active Leasequery with Topology Discovery

Enable active leasequery with topology discovery on the pair of DHCP relay agents that support a given subscriber redundancy group. You must repeat the configuration for each pair of relay agents for different redundancy groups.



NOTE: The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level.



NOTE: For dual-stack subscribers, you must configure active leasequery with topology discovery for both DHCPv4 and DHCPv6.



NOTE: Because active leasequery is an extension of bulk leasequery, you must also configure bulk leasequery for active leasequery to operate. You must configure bulk leasequery before you configure active leasequery. See ["Configuring and Using DHCP Bulk Leasequery" on page 1141](#).

1. Specify that you want to configure active leasequery options for the DHCP relay agent.

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

2. Specify the IP address for a peer with which this relay agent synchronizes information. You must also configure active leasequery on the peer.



NOTE: This is the address used for the TCP connection. It can be a physical interface address or a loopback address per pair of peers.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set peer-address ip-address
```

3. Configure the relay agent to send topology discovery messages to determine the remote access interfaces for subscriber redundancy groups on similarly configured peer relay agents. Discovering

the topology enables the relay agents to build translation tables of local and remote interfaces to support an interface-level, primary/backup redundancy scheme.

```
[edit forwarding-options dhcp-relay active-leasequery]
user@host# set topology-discover
```

4. Configure the relay agent to always include Option 82, suboption 1, the Agent Circuit ID. This is the name of the access interface.



NOTE: For DHCPv6, the equivalent statement is *relay-agent-interface-id* to include Option 18.

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

Verifying M:N Redundancy and Active Leasequery Topology Discovery Information

IN THIS SECTION

- Purpose | 1554
- Action | 1554

Purpose

Determine status information and statistics for access interfaces, relay agents, and subscribers that are part of your topology for M:N redundancy with DHCP active leasequery topology discovery.

Action

- To verify the VRRP redundancy state of access interfaces:

```
user@host>show vrrp
```

- To verify that the redundancy state of a specified access logical interface is Master on the primary relay agent and Backup on the backup relay agent:

```
user@host>show system subscriber-management redundancy-state dhcp active-leasequery interface
interface-name
```

This interface can be either a subscriber interface or the underlying VLAN interface. For VRRP redundancy, the redundancy state is the same as the VRRP state of the underlying logical interface. For pseudowire redundancy, the redundancy state is based on the state of the pseudowire interface.

- To verify that subscribers in a redundancy group are active on both the primary and backup relay agents:

```
user@host>show subscribers option
```

The `show subscribers` command has a number of options; you can display subscribers by IP address, interface name, VLAN ID, Agent Circuit ID, subscriber state, and so on.

- To verify that the DHCP relay binding information is the same for the subscribers in a redundancy group on both the primary and backup relay agents:

```
user@host>show dhcp relay binding verbose
user@host>show dhcpv6 relay binding verbose
```

You can also specify results for an IP address or an interface.

- To view a list of all active leasequery peers:

```
user@host>show dhcp relay active-leasequery summary
user@host>show dhcpv6 relay active-leasequery summary
```

- To view the topology discovery translation table for a peer relay agent, including the local and remote circuit IDs (access interfaces), local access interface address, transaction ID (xid), and the state of topology discovery, redundancy, and subscriber synchronization:

```
user@host>show dhcp relay active-leasequery peer address details
user@host>show dhcpv6 relay active-leasequery peer address details
```

- To view active leasequery statistics, such as the number of DHCP bindings sent or received for an interface or peer.

```

user@host>show dhcp relay active-leasequery statistics (interface interface-name | peer ip-address)
user@host>show dhcpv6 relay active-leasequery statistics (interface interface-name | peer ipv6-address)

```

- To clear active leasequery statistics.

```

user@host>clear dhcp relay active-leasequery statistics (interface interface-name | peer ip-address)
user@host>clear dhcpv6 relay active-leasequery statistics (interface interface-name | peer ipv6-address)

```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, you can use pseudowire redundancy to provide M:N redundancy when the access network consists of Layer 2 (L2) circuits over IP/MPLS.
19.2R1	Starting in Junos OS Release 19.2R1, you can configure M:N subscriber redundancy as a mechanism for improving network resiliency by protecting subscribers from a variety of software and hardware failures.

RELATED DOCUMENTATION

Understanding VRRP
DHCP Leasequery Methods 1116
DHCP Active Leasequery 1127

M:N Subscriber Service Redundancy on DHCP Server

SUMMARY

Learn about M:N subscriber redundancy on DHCP server, which ensures uninterrupted subscriber service.

IN THIS SECTION

- [M:N Subscriber Service Redundancy on DHCP Server Overview | 1557](#)

M:N Subscriber Service Redundancy on DHCP Server Overview

IN THIS SECTION

- [Benefits of M:N Subscriber Service Redundancy on DHCP Server | 1561](#)

You can configure M:N subscriber service redundancy on DHCP server running on MX Series broadband network gateway (BNG). DHCP server maintains considerable amount of authoritative information regarding the address it has leased to the DHCP clients. To achieve MX Series chassis level BNG redundancy for broadband subscribers, the backup MX Series device running DHCP server should possess all the subscriber authoritative information. The backup server ensures uninterrupted subscriber service when you reboot or replace the primary DHCP server, or the primary server has any hardware failures such as access link failures, access line card failure, or chassis failure.

Subscriber service redundancy on DHCP server focuses on subscriber synchronization between the peer servers using active leasequery. Live update of binding information between the two peer servers help to maintain the servers in hot standby mode.

In 1:1 subscriber service redundancy, one DHCP server (primary DHCP server) is backed up on another DHCP server (backup DHCP server) for subscriber interfaces. Both the primary and the backup servers have the same set of configurations. Starting Junos OS Release 24.4R1, chassis based DHCP redundancy feature for MX480 access models, provides the additional functionality of 1:1 chassis based active lease query, below limit of quantification (ALQ/BLQ) redundancy for non-participating underlying subscriber interfaces without the need to configure topology-discovery. This feature is supported on stacks and DHCP configurations, BBE and non-BBE DHCP without topology discovery, for the following scenarios:

- Subscriber management enabled.
- Subscriber management disabled.

- IP Demux/ IP Demux Lite.
- Dual Stack and Dual Stack single session.
- Pseudo Wire Access model PS Interface (L2 Circuit / EVPN VPWS / L2VPN).
- VRRP Access model GE, XE, and AE Interfaces.
- Non-default routing instance.
- DHCP Relay and DHCP Server.

See Configuration Example: ALQ/BLQ for 1:1 DHCP Redundancy for details on configuration.

In M:N subscriber service redundancy multiple (M) DHCP servers (primary DHCP server) are backed up on multiple (N) DHCP servers (backup DHCP server). The M:N subscriber service redundancy requires topology discovery to map the interfaces between the peer servers. To replicate the subscribers on interface, the active leasequery uses Gi-Address query for IPv4 and link-address query for IPv6.

When the subscribers receive the leasequery response, the relevant state machine power up the subscriber in the backup server. Then the DHCP address and lease information synchronizes between the servers. If the lease or address information changes, the backup BNG runs through the relevant state machine to power up or down the subscriber state.

Currently the subscriber services redundancy on DHCP supports pseudowire redundancy protocol and topology discovery over pseudowire between the peer servers. The subscriber services redundancy supports the protocols listed in [Table 63 on page 1558](#).

Table 63: Subscriber Services Redundancy

Supported Protocols	Subscriber Services Redundancy Mode	Additional Details
IPoE DHCP relay, static VLAN	M:N stateful with VRRP and active leasequery	For dynamic VLAN support must use PWHT
IPoE DHCP relay over PWHT	M:N stateful with active leasequery	
IPoE DHCP server over PWHT	M:N stateful	Include dynamic or static VLAN support

[Figure 42 on page 1559](#) shows the topology for L2 circuit based IP/MPLS PWHT in client-server mode.

Figure 42: L2 Circuit Based IP/MPLS PWHT in Client-Server Mode

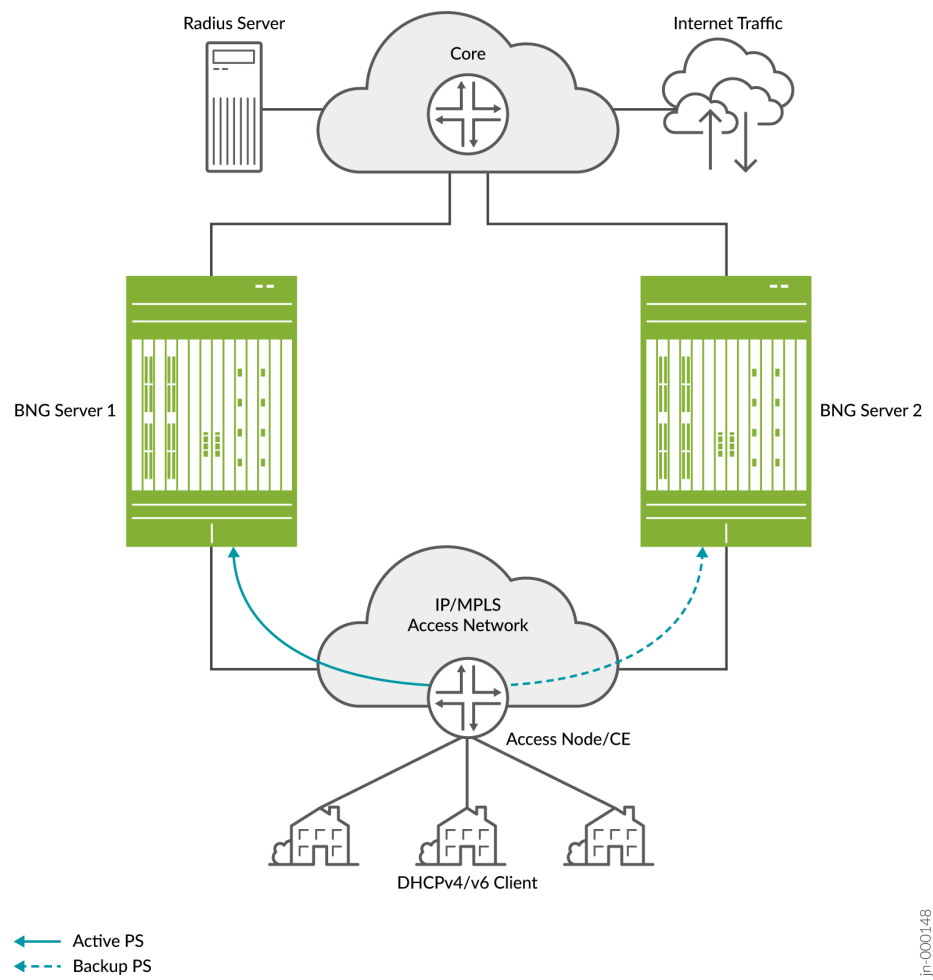



Figure 43 on page 1560 shows the topology for L2 circuit based IP/MPLS PWHT in client-relay-server mode.

Figure 43: L2 Circuit Based IP/MPLS PWHT in Client-Relay-Server Mode



In both client-server mode and client-relay-server mode topology, BNG servers use TCP connection for active leasequery to synchronize binding details. The subscriber service redundancy on DHCP server occurs in the following order:

1. Active pseudowire link receives packets from the client.
2. Subscriber connects to the primary BNG.
3. Primary BNG synchronizes the subscriber binding details to backup BNG using TCP connection.
4. When you reboot or replace the primary BNG or the primary BNG has any chassis failure the backup pseudowire link become active.
5. The backup BNG receives packets from the client.
6. As backup BNG was already in hot-standby mode it can renew or rebind packets for active leasequery and synchronize subscribers as well.

For M:N subscriber service redundancy, you need to backup the subscribers interface on the backup DHCP server. The interface can have different names. The primary DHCP server uses the topology discovery to map the interfaces between peer DHCP servers.

The DHCP server uses the Gi-address or link-address query to replicate the subscribers information on the backup DHCP server. In server, clients having different Gi-address or link-address comes up on single interface, thus the primary BNG should respond the query with all the subscribers having different Gi-addresses or link-address on interface. To support this functionality, the server creates a new table to store the clients based on incoming interface. When the server receives a Gi-address or link-address query, the server responds the query from the new table as follows:

- When the server sends request, it checks the topology discovery configuration and sends GI-address or link-address based query with interface IP address.
- When the server receives a GI-address or link-address based query, server checks the existing server configuration. If an active leasequery configuration is available, the server responds to the query based the new database.

Active leasequery can be done between the relay to relay or server to server at any time. DHCP server may not accept connection from peer server or relay simultaneously, thus the configuration in DHCP server can be either of active-leasequery, or allow-active-leasequery, allow-bulk-leasequery, or allow-leasequery.

Benefits of M:N Subscriber Service Redundancy on DHCP Server

- Provides uninterrupted subscriber services at DHCP server level.

SEE ALSO

<https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/dhcp-local-server-active-leasequery.html>

N+1 Support for BNG M:N Subscriber Service Redundancy

SUMMARY

Learn about N+1 support for broadband network gateway (BNG) M:N subscriber service redundancy, which provides remarkable reduction in the reserved resources for the backup BNG.

IN THIS SECTION

- [N+1 Support for BNG M:N Subscriber Service Redundancy Overview | 1562](#)
- [How N+1 Support for BNG M:N Subscriber Service Redundancy Works | 1562](#)

N+1 Support for BNG M:N Subscriber Service Redundancy Overview

The N+1 support for BNG M:N subscriber service redundancy is a mechanism to back up multiple primary BNGs to a single backup BNG. This mechanism provides reduction in the reserved resources for redundancy purpose by over-subscribing the secondary Packet Forwarding Engine in backup chassis.

How N+1 Support for BNG M:N Subscriber Service Redundancy Works

IN THIS SECTION


- [Subscriber Service Redundancy When Primary BNG Fails | 1564](#)
- [Subscriber Service Revert When the Primary BNG Becomes Active | 1564](#)

[Figure 44 on page 1563](#) illustrates N+1 support for BNG M:N subscriber service redundancy. There are four BNGs shown in the topology. The BNGs A, C, and D are the active BNGs with 64000 dual-stack subscribers on each BNG. The backup BNG B with one line card backing up the other three active BNGs. You can use any MX Series device that can support MPC7 or MX10003 device with LC2103 as a backup BNG.

A1, C1, and D1 are the primary subscriber redundancy groups handling traffic of 64000 subscribers on each BNG. A2, C2, and D2 are the secondary subscriber redundancy groups in service-activation-on-failover mode.

By default, the M:N subscriber redundancy feature configures the backup BNG in hot-standby mode. To specifically enable the Packet Forwarding Engine over-subscription, you need to configure the service-activation-on-failover mode on the backup BNG.

Figure 44: N+1 Support for BNG M:N Subscriber Service Redundancy



In this redundancy model we've introduced a service-activation-on-failover mode. In the service-activation-on-failover mode, you can configure the subscriber state for an interface using less resources in the backup BNG to forward traffic. When the primary BNG fails, the traffic switches over to the backup BNG with basic statistics. The additional services such as CoS and firewall automatically come to action in the background after the backup interface becomes active and consumes the additional resources. The operational state of the backup interface transitions from basic forwarding to full service restoration.

The new programming mode enables the system to consume less resources on the backup BNG. Hence, you can back up more subscribers when the Packet Forwarding Engine is not handling any traffic. This backup subscription is known as Packet Forwarding Engine over-subscription on the backup BNG. With the service-activation-on-failover mode, you can host three times more subscribers on the backup BNG than the primary BNGs. The n+1 support for BNG mode reduces the cost of deploying backup BNGs.

When the subscribers log in to the primary BNGs, the active leasequery brings the subscriber state to the backup BNG. As the backup BNG hosts the service-activation-on-failover mode, the backup BNG consumes minimal Packet Forwarding Engine resources and back ups up to 192000 subscribers.

Subscriber Service Redundancy When Primary BNG Fails

Let's see how the system manages when a BNG fails or a BNG becomes inactive. Considering the [Figure 44 on page 1563](#), when the BNG C fails, the subscribers connected to the BNG C re-routes the traffic through the backup BNG B. As soon as the traffic re-routes to the secondary subscriber redundancy group C2, the BNG B performs the following:

- Starts forwarding the upstream and downstream traffic immediately with best-effort.
- Initiates background programming for the services such as CoS and firewall by utilizing the additional resources allocated in BNG B.
- The BNG B restores the full SLA for subscribers and the operational state becomes full-service when the background programming completes.
- The other secondary subscriber redundancy groups A2 and D2 continue to back up the BNGs A and D.

Subscriber Service Revert When the Primary BNG Becomes Active

You can configure the primary BNG C to revert the traffic flow from the backup BNG to the primary BNG when it becomes active. We recommend to use manual revert after checking the state of both BNGs for subscriber programming and confirming that a revert back will succeed. Consider the following scenarios when you enable the auto-revert traffic switchover functionality:

- If the primary BNG fails due to link failure, the background programming of the backup BNG takes several minutes depending on the number of subscribers. A quick revert is not desirable.
- If the primary BNG fails due to the line card or chassis failure, the time to synchronize the original primary chassis or line card using active leasequery or bulk leasequery depends on the number of subscribers.
- The system requires more time to analyze unplanned failures and make the line card or chassis into active service.

N+1 support for BNG M:N subscriber service redundancy does not support redundancy on multiple BNG failures at a time. If multiple BNGs fail at a time, the system back ups only the first BNG. The data of the remaining subscribers on the other failed BNGs are lost completely.

RELATED DOCUMENTATION

show system subscriber-management redundancy-state interface

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery

SUMMARY

Learn about broadband network gateway (BNG) redundancy using packet triggered based recovery which provides simple, easy to use, and lightweight stateless subscriber redundancy.

IN THIS SECTION

- [BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Overview | 1565](#)
- [How BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Works | 1566](#)
- [Configuring BNG Redundancy Using Packet Triggered Recovery for DHCP Subscribers | 1568](#)

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Overview

The BNG redundancy for DHCP subscribers using packet triggered based recovery provides simple, easy to use, and lightweight stateless redundancy with minimal traffic loss. The stateless BNG redundancy for DHCP subscribers supports dynamic C-VLAN and static VLAN model for both relay and server. The packet triggered based recovery utilizes the existing features such as auto configuration of VLAN and packet triggered subscribers.

Auto Configuration of VLAN

The auto configuration feature creates dynamic VLAN (DVLAN) logical interface on receiving the first VLAN packet from the client. On receiving the first packet, the Routing Engine authenticates the subscriber with authenticating server. The authentication server might need the accounting and advanced services details for authenticating the subscriber. The Routing Engine creates the DVLAN logical interface based on the request from the authenticating server. After creating the DVLAN logical interface, the system forwards the packet to the protocol stack for further processing.

Packet Triggered Subscribers

The packet-triggered subscriber feature creates IP demux logical interface on receiving a packet from clients with the pre-assigned IPv4 or IPv6 address. The forwarding plane validates the source IP address and matches with the configured IP address or prefix ranges. After the source IP address validation, the

forwarding plane forwards the packet to the Routing Engine. The Routing Engine authenticates the subscriber with authenticating server as per the volume of accounting and advanced services such as firewall filter and CoS. Routing Engine creates IP demux logical interface as per the services requested by the authenticating server.

Benefits of BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery

- Provides simple backup BNG deployment.

How BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Works

Primary BNG hosts the subscribers during normal traffic flow. When the traffic flow fails in the primary BNG, the access nodes redirect the traffic to the backup BNG. The primary BNG can fail due to following reasons:

- Intermediate node failure or link failure which breaks the MPLS path between access node and primary BNG.
- Primary BNG link or port failure.
- Primary BNG line card failure.
- Primary BNG Routing Engine failure.
- Primary BNG chassis failure.
- Primary BNG to core network link failure.

Figure 45: L2 Circuit Based on IP/MPLS PWHT Scenario




Figure 45 on page 1567 shows the topology diagram for layer 2 circuit based on IP/MPLS pseudowire headend termination (PWHT) scenario.

Based on the first traffic after failover, the Routing Engine creates DVLAN and dynamic IP subscriber. Packet Forwarding Engine forwards the subsequent traffic in the forwarding plane to the core router as per QoS and services attached to the IP subscriber. This QoS and the services are not the same QoS and services of the subscriber created in the Primary BNG. These are the common default dynamic IP subscriber profile features, assigned by RADIUS server or local configuration, until the session lease renewal and re-authentication occurs.

Once the system creates the DHCP subscriber in the secondary BNG, it provides limited QoS and other services at best-effort traffic with minimal interruption. When the DHCP client lease timer expires, it tries to re-negotiate lease time and a new DHCP protocol exchange takes place. This time, the system creates the fully functional DHCP subscriber along with QoS and advanced services as that of the primary BNG. The Packet Forwarding Engine forwards the traffic also to the core router accordingly. The system deletes the dynamic IP subscriber when the fully functional DHCP subscriber is active.

The traffic switchover to the backup BNG and the revert to the primary BNG process is similar. If revert occurs after the first lease timeout, the system proceeds with the switchover process. If revert occurs before the first lease timeout, the system proceeds with revert as it still has the previously assigned IP address and DHCP bindings.

The BNG redundancy using packet triggered based recovery feature supports the following access network topology for BNG redundancy:

- Layer 2 VPN scenario
- Layer 2 circuit based on IP/MPLS PWHT scenario
- Ethernet VPN–virtual private wireless service (EVPN-VPWS).

Configuring BNG Redundancy Using Packet Triggered Recovery for DHCP Subscribers

IN THIS SECTION

- [Overview | 1568](#)
- [Requirements | 1568](#)
- [Topology | 1569](#)
- [Configuration | 1570](#)
- [Verification | 1614](#)

Overview

Starting from 22.4R1, Junos supports Broadband Network Gateway (BNG) redundancy configuration using packet-trigger based recovery that provides easy to use and lightweight stateless subscriber redundancy.

This section provides a configuration example for packet triggered BNG redundancy for DHCP subscribers using a local DHCP server.

Requirements

The configuration example uses the following devices:

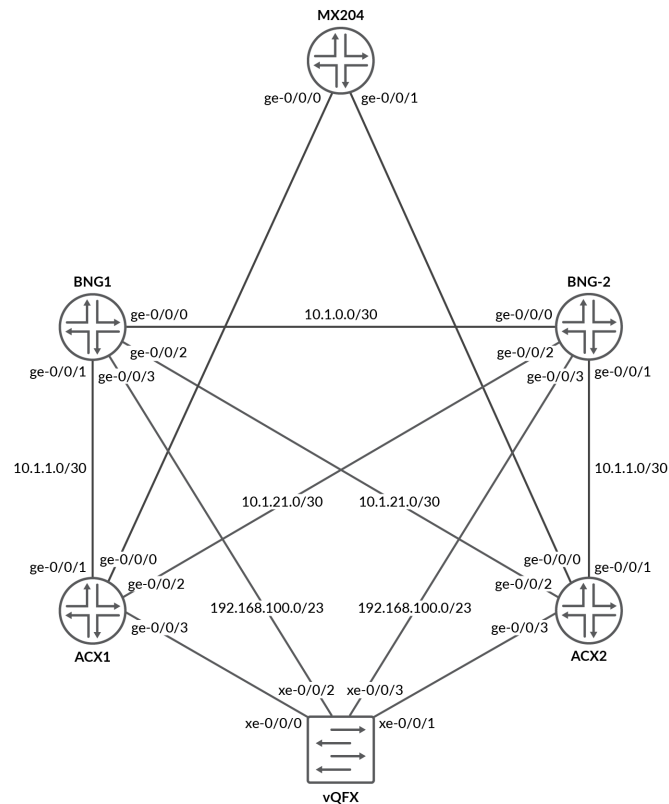
- BNG1 and BNG2 broadband network gateways run EVPN-VPWS pseudowire headend termination (PWHT) with the ACX aggregation nodes and terminate the IPoE sessions. The BNG implements the packet triggered redundancy for the IPoE sessions.

- ACX1 and ACX2 devices aggregate the Access Nodes traffic through the Cloud Metro Fabric towards the BNGs.
- MX204 device is used for simulated peripheral connectivity.
- vQFX instance for Q-in-Q tunneling and VLAN translation.

Topology

Figure 2 on page 1569 shows the physical topology with two vMX devices configured as BNG1 and BNG2 servers, two access devices ACX1 and ACX2, a vQFX and an MX204 device.

Figure 46: Topology



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1570](#)
- [Device BNG1 | 1570](#)
- [Device BNG2 | 1578](#)
- [Device ACX1 | 1585](#)
- [Device ACX2 | 1587](#)
- [Step-by-Step Procedure | 1589](#)
- [Configuring BNG1 | 1589](#)
- [Configuring BNG2 | 1598](#)
- [Configuring ACX1 | 1608](#)
- [Configuring ACX2 | 1611](#)

CLI Quick Configuration

Device BNG1

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

```
set system host-name BNG1
set system configuration-database max-db-size 698343424
set system services ssh root-login allow
set chassis network-services enhanced-ip
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED *****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
```

```

set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/0 apply-groups GR-CORE-INTF
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.21.2/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 mtu 9192
set interfaces ge-0/0/3 unit 0 family inet address 192.168.100.171/23
set interfaces lo0 unit 0 family inet address 172.31.100.3/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0003.00
set interfaces lo0 unit 0 family mpls
set protocols isis interface ge-0/0/0.0
set protocols isis interface ge-0/0/1.0
set protocols isis interface ge-0/0/2.0
set protocols isis interface lo0.0 passive
set protocols isis level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
set routing-options router-id 172.31.100.3
set routing-options autonomous-system 65000

```

```

set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.3
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP cluster 1.1.1.1
set protocols bgp group IBGP neighbor 172.31.100.4
set protocols bgp group IBGP neighbor 172.31.100.11
set protocols bgp group IBGP neighbor 172.31.100.12
set protocols ldp deaggregate
set protocols ldp transport-address 172.31.100.3
set protocols ldp interface all
set protocols mpls interface all
set system services subscriber-management traceoptions file submgmt.log
set system services subscriber-management traceoptions file size 30m
set system services subscriber-management traceoptions file files 10
set system services subscriber-management traceoptions flag all
set system services subscriber-management gres-route-flush-delay
set system services subscriber-management enable
set system processes general-authentication-service traceoptions file authd
set system processes general-authentication-service traceoptions file size 10m
set system processes general-authentication-service traceoptions file files 10
set system processes general-authentication-service traceoptions flag all
set system processes smg-service traceoptions file smgd
set system processes smg-service traceoptions file size 10m
set system processes smg-service traceoptions file files 10
set system processes smg-service traceoptions level all
set system processes smg-service traceoptions flag all
set system processes dhcp-service traceoptions file jdhcpd
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions file files 10
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set system processes dhcp-service traceoptions flag all
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis fpc 0 performance-mode
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
set chassis pseudowire-service device-count 10
set interfaces ps0 anchor-point lt-0/0/10
set interfaces ps0 flexible-vlan-tagging
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include delimiter
"@"
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include user-

```

```

prefix vlan
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include interface-
name
set interfaces ps0 auto-configure stacked-vlan-ranges access-profile no-auth
set interfaces ps0 auto-configure remove-when-no-subscribers
set interfaces ps0 mtu 2022
set interfaces ps0 esi 00:10:00:00:01:00:00:00:10:00
set interfaces ps0 esi single-active
set interfaces ps0 esi df-election-type preference value 1000
set interfaces ps0 unit 0 encapsulation ethernet-ccc
set interfaces lo0 unit 20 description "VRF:internet Loopback"
set interfaces lo0 unit 20 family inet address 172.16.100.3/32 primary
set routing-instances internet instance-type vrf
set routing-instances internet routing-options auto-export
set routing-instances internet interface lo0.20
set routing-instances internet route-distinguisher 172.31.100.3:12
set routing-instances internet vrf-import internet-vrf-import-pol
set routing-instances internet vrf-export internet-vrf-export-pol
set routing-instances internet vrf-table-label
set policy-options policy-statement internet-vrf-export-pol term all then community add 65000:999
set policy-options policy-statement internet-vrf-export-pol term all then accept
set policy-options policy-statement internet-vrf-import-pol term default from community 65000:999
set policy-options policy-statement internet-vrf-import-pol term default then accept
set policy-options policy-statement internet-vrf-import-pol term subs from community 65000:1131
set policy-options policy-statement internet-vrf-import-pol term subs then accept
set policy-options policy-statement internet-vrf-import-pol term other from community 65000:113
set policy-options policy-statement internet-vrf-import-pol term other then accept
set policy-options community 65000:113 members target:65000:113
set policy-options community 65000:1131 members target:65000:1131
set policy-options community 65000:999 members target:65000:999
set interfaces lo0 unit 313 description "VRF:dhcp-subs Loopback"
set interfaces lo0 unit 313 family inet address 172.16.16.3/32
set interfaces lo0 unit 313 family inet address 10.42.0.1/32 primary
set interfaces lo0 unit 313 family inet6 address 2015:cafe:2000::1/128
set routing-instances dhcp-subs instance-type vrf
set routing-instances dhcp-subs routing-options rib dhcp-subs.inet6.0 static route ::/0 next-
table internet.inet6.0
set routing-instances dhcp-subs routing-options rib dhcp-subs.inet6.0 static route ::/0 no-
readvertise
set routing-instances dhcp-subs routing-options router-id 172.16.16.3
set routing-instances dhcp-subs routing-options flow term-order standard
set routing-instances dhcp-subs routing-options static route 0.0.0.0/0 next-table internet.inet.0
set routing-instances dhcp-subs routing-options static route 0.0.0.0/0 no-readvertise

```

```

set routing-instances dhcp-subs routing-options static route 10.42.0.0/16 discard
set routing-instances dhcp-subs routing-options static route 10.42.0.0/16 tag 200
set routing-instances dhcp-subs routing-options auto-export
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls liveness-
detection failure-action clear-binding-if-interface-up
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls liveness-
detection method bfd version automatic
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls liveness-
detection method bfd minimum-interval 30000
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls liveness-
detection method bfd multiplier 3
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls overrides client-
discover-match incoming-interface
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls overrides dual-
stack dhcp-ds
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls interface
demux0.0
set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls interface ps0.0
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
authentication username-include mac-address
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds on-
demand-address-allocation
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
classification-key mac-address
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
protocol-master inet
set routing-instances dhcp-subs system services dhcp-local-server no-stale-timer-refresh
set routing-instances dhcp-subs system services dhcp-local-server stale-timer 60
set routing-instances dhcp-subs access address-assignment high-utilization 80
set routing-instances dhcp-subs access address-assignment abated-utilization 70
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet network 10.42.0.0/16
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet range range1 low 10.42.0.2
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet range range1 high 10.42.255.254
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes maximum-lease-time 600
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes server-identifier 10.42.0.1
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes router 10.42.0.1
set routing-instances dhcp-subs access-profile no-auth

```

```

set routing-instances dhcp-sub interface lo0.313
set routing-instances dhcp-sub route-distinguisher 172.31.100.3:13
set routing-instances dhcp-sub vrf-import dhcp-sub-vrf-import-pol
set routing-instances dhcp-sub vrf-export dhcp-sub-vrf-export-pol
set routing-instances dhcp-sub vrf-table-label
set policy-options policy-statement dhcp-sub-vrf-export-pol term loop from protocol direct
set policy-options policy-statement dhcp-sub-vrf-export-pol term loop from route-filter
172.16.16.3/32 exact
set policy-options policy-statement dhcp-sub-vrf-export-pol term loop then community add
65000:113
set policy-options policy-statement dhcp-sub-vrf-export-pol term pools from protocol static
set policy-options policy-statement dhcp-sub-vrf-export-pol term pools from tag 200
set policy-options policy-statement dhcp-sub-vrf-export-pol term pools then community add
65000:113
set policy-options policy-statement dhcp-sub-vrf-export-pol term subs then community add
65000:1131
set policy-options policy-statement dhcp-sub-vrf-export-pol term subs then accept
set policy-options policy-statement dhcp-sub-vrf-import-pol term all from community 65000:111
set policy-options policy-statement dhcp-sub-vrf-import-pol term all then accept
set policy-options community 65000:111 members target:65000:111
set policy-options community 65000:111 members target:65000:11
set access domain map none access-profile no-auth
set access domain map none target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set access domain map ps0:* access-profile no-auth
set access domain map ps0:* target-routing-instance dhcp-sub
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-address
255.255.255.255/32
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-port dhcp
set firewall family inet filter rpf-pass-dhcp term allow-dhcp then accept
set firewall family inet filter rpf-pass-dhcp term default then discard
set system dynamic-profile-options versioning
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" actual-
transit-statistics
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" demux-options
underlying-interface "$junos-underlying-interface"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
rpf-check fail-filter rpf-pass-dhcp

```

```

set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
demux-source $junos-subscriber-ip-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
unnumbered-address lo0.313
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet6
demux-source $junos-subscriber-ipv6-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet6
unnumbered-address lo0.313
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
dynamic-profile prod-dhcp-base
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" interface "$junos-
interface-name"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix next-hop "$junos-framed-route-nexthop"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles auto-pwht interfaces interface-set "$junos-phy-ifd-interface-set-name"
interface "$junos-interface-ifd-name" unit "$junos-interface-unit"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" no-traps
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" proxy-arp restricted
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" vlan-id "$junos-vlan-id"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" family inet mac-validate loose
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags outer "$junos-stacked-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags inner "$junos-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate strict
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-

```

```

interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile PROF_AUTOSENSE_IPDEMUX
network 10.42.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile PROF_AUTOSENSE_IPDEMUX
network 10.43.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
delimiter "@"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include user-
prefix vlan
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
interface-name
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges session-timeout 600
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-
unit" family inet mac-validate strict
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-
unit" family inet unnumbered-address lo0.313
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-
unit" family inet6 unnumbered-address lo0.313
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht accept
any
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht ranges
any,any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht accept any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht ranges any
set routing-instances EVPN-VPWS-BNG-1 instance-type evpn-vpws
set routing-instances EVPN-VPWS-BNG-1 protocols evpn interface ps0.0 vpws-service-id local 9999
set routing-instances EVPN-VPWS-BNG-1 protocols evpn interface ps0.0 vpws-service-id remote 1111
set routing-instances EVPN-VPWS-BNG-1 interface ps0.0
set routing-instances EVPN-VPWS-BNG-1 route-distinguisher 172.31.100.3:11
set routing-instances EVPN-VPWS-BNG-1 vrf-target target:65000:11

```


Device BNG2

```

set system host-name BNG2
set system configuration-database max-db-size 698343424
set system services ssh root-login allow
set chassis network-services enhanced-ip
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-ADD-
DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/0 apply-groups GR-CORE-INTF
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.2/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.12.2/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 mtu 9192
set interfaces ge-0/0/3 unit 0 family inet address 192.168.100.171/23
set interfaces lo0 unit 0 family inet address 172.31.100.4/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0004.00
set interfaces lo0 unit 0 family mpls
set protocols isis interface ge-0/0/0.0

```

```

set protocols isis interface ge-0/0/1.0
set protocols isis interface ge-0/0/2.0
set protocols isis interface lo0.0 passive
set protocols isis level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
set routing-options router-id 172.31.100.4
set routing-options autonomous-system 65000
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.4
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP cluster 2.2.2.2
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.11
set protocols bgp group IBGP neighbor 172.31.100.12
set protocols ldp deaggregate
set protocols ldp transport-address 172.31.100.4
set protocols ldp interface all
set protocols mpls interface all
set system services subscriber-management traceoptions file submgmt.log
set system services subscriber-management traceoptions file size 30m
set system services subscriber-management traceoptions file files 10
set system services subscriber-management traceoptions flag all
set system services subscriber-management gres-route-flush-delay
set system services subscriber-management enable
set system processes general-authentication-service traceoptions file authd
set system processes general-authentication-service traceoptions file size 10m
set system processes general-authentication-service traceoptions file files 10
set system processes general-authentication-service traceoptions flag all
set system processes smg-service traceoptions file smgd
set system processes smg-service traceoptions file size 10m
set system processes smg-service traceoptions file files 10
set system processes smg-service traceoptions level all
set system processes smg-service traceoptions flag all
set system processes dhcp-service traceoptions file jdhcpd

```

```

set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions file files 10
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set system processes dhcp-service traceoptions flag all
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis fpc 0 performance-mode
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
set chassis pseudowire-service device-count 10
set interfaces ps0 anchor-point lt-0/0/10
set interfaces ps0 flexible-vlan-tagging
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include delimiter
"@"
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include user-
prefix vlan
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include interface-
name
set interfaces ps0 auto-configure stacked-vlan-ranges access-profile no-auth
set interfaces ps0 auto-configure remove-when-no-subscribers
set interfaces ps0 mtu 2022
set interfaces ps0 esi 00:10:00:00:01:00:00:00:10:00
set interfaces ps0 esi single-active
set interfaces ps0 esi df-election-type preference value 999
set interfaces ps0 unit 0 encapsulation ethernet-ccc
set interfaces lo0 unit 20 description "VRF:internet Loopback"
set interfaces lo0 unit 20 family inet address 172.16.100.4/32 primary
set routing-instances internet instance-type vrf
set routing-instances internet routing-options auto-export
set routing-instances internet interface lo0.20
set routing-instances internet route-distinguisher 172.31.100.4:12
set routing-instances internet vrf-import internet-vrf-import-pol
set routing-instances internet vrf-export internet-vrf-export-pol
set routing-instances internet vrf-table-label
set policy-options policy-statement internet-vrf-export-pol term all then community add 65000:999
set policy-options policy-statement internet-vrf-export-pol term all then accept
set policy-options policy-statement internet-vrf-import-pol term default from community 65000:999
set policy-options policy-statement internet-vrf-import-pol term default then accept
set policy-options policy-statement internet-vrf-import-pol term subs from community 65000:1131
set policy-options policy-statement internet-vrf-import-pol term subs then accept
set policy-options policy-statement internet-vrf-import-pol term other from community 65000:113
set policy-options policy-statement internet-vrf-import-pol term other then accept

```

```

set policy-options community 65000:113 members target:65000:113
set policy-options community 65000:1131 members target:65000:1131
set policy-options community 65000:999 members target:65000:999
set interfaces lo0 unit 313 description "VRF:dhcp-sub's Loopback"
set interfaces lo0 unit 313 family inet address 172.16.16.4/32
set interfaces lo0 unit 313 family inet address 10.43.0.1/32 primary
set interfaces lo0 unit 313 family inet6 address 2016:cafe:2000::1/128
set routing-instances dhcp-sub's instance-type vrf
set routing-instances dhcp-sub's routing-options rib dhcp-sub's.inet6.0 static route ::/0 next-
table internet.inet6.0
set routing-instances dhcp-sub's routing-options rib dhcp-sub's.inet6.0 static route ::/0 no-
readvertise
set routing-instances dhcp-sub's routing-options router-id 172.16.16.4
set routing-instances dhcp-sub's routing-options flow term-order standard
set routing-instances dhcp-sub's routing-options static route 0.0.0.0/0 next-table internet.inet.0
set routing-instances dhcp-sub's routing-options static route 0.0.0.0/0 no-readvertise
set routing-instances dhcp-sub's routing-options static route 10.43.0.0/16 discard
set routing-instances dhcp-sub's routing-options static route 10.43.0.0/16 tag 200
set routing-instances dhcp-sub's routing-options auto-export
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls liveness-
detection failure-action clear-binding-if-interface-up
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls liveness-
detection method bfd version automatic
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls liveness-
detection method bfd minimum-interval 30000
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls liveness-
detection method bfd multiplier 3
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls overrides client-
discover-match incoming-interface
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls overrides dual-
stack dhcp-ds
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls interface
demux0.0
set routing-instances dhcp-sub's system services dhcp-local-server group dhcp-ls interface ps0.0
set routing-instances dhcp-sub's system services dhcp-local-server dual-stack-group dhcp-ds
authentication username-include mac-address
set routing-instances dhcp-sub's system services dhcp-local-server dual-stack-group dhcp-ds on-
demand-address-allocation
set routing-instances dhcp-sub's system services dhcp-local-server dual-stack-group dhcp-ds
classification-key mac-address
set routing-instances dhcp-sub's system services dhcp-local-server dual-stack-group dhcp-ds
protocol-master inet
set routing-instances dhcp-sub's system services dhcp-local-server no-stale-timer-refresh

```

```

set routing-instances dhcp-subsystem services dhcp-local-server stale-timer 60
set routing-instances dhcp-subsystem access address-assignment high-utilization 80
set routing-instances dhcp-subsystem access address-assignment abated-utilization 70
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet network 10.43.0.0/16
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet range range1 low 10.43.0.2
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet range range1 high 10.43.254.254
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes maximum-lease-time 600
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes server-identifier 10.43.0.1
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0 family
inet dhcp-attributes router 10.43.0.1
set routing-instances dhcp-subsystem access-profile no-auth
set routing-instances dhcp-subsystem interface lo0.313
set routing-instances dhcp-subsystem route-distinguisher 172.31.100.4:13
set routing-instances dhcp-subsystem vrf-import dhcp-subsystem-vrf-import-pol
set routing-instances dhcp-subsystem vrf-export dhcp-subsystem-vrf-export-pol
set routing-instances dhcp-subsystem vrf-table-label
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop from protocol direct
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop from route-filter
172.16.16.4/32 exact
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop then community add
65000:113
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools from protocol static
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools from tag 200
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools then community add
65000:113
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term subs then community add
65000:1131
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term subs then accept
set policy-options policy-statement dhcp-subsystem-vrf-import-pol term all from community 65000:111
set policy-options policy-statement dhcp-subsystem-vrf-import-pol term all then accept
set policy-options community 65000:111 members target:65000:111
set policy-options community 65000:111 members target:65000:11
set policy-options community 65000:113 members target:65000:113
set policy-options community 65000:1131 members target:65000:1131
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
set access domain map none access-profile no-auth

```

```

set access domain map none target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set access domain map ps0:* access-profile no-auth
set access domain map ps0:* target-routing-instance dhcp-sub
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-address
255.255.255.255/32
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-port dhcp
set firewall family inet filter rpf-pass-dhcp term allow-dhcp then accept
set firewall family inet filter rpf-pass-dhcp term default then discard
set system dynamic-profile-options versioning
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" actual-
transit-statistics
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" demux-options
underlying-interface "$junos-underlying-interface"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
rpf-check fail-filter rpf-pass-dhcp
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
demux-source $junos-subscriber-ip-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet
unnumbered-address lo0.313
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet6
demux-source $junos-subscriber-ipv6-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family inet6
unnumbered-address lo0.313
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds
dynamic-profile prod-dhcp-base
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" interface "$junos-
interface-name"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix next-hop "$junos-framed-route-nexthop"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles auto-pwht interfaces interface-set "$junos-phy-ifd-interface-set-name"
interface "$junos-interface-ifd-name" unit "$junos-interface-unit"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" no-traps
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" proxy-arp restricted

```

```

set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" vlan-id "$junos-vlan-id"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" family inet mac-validate loose
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-interface-
unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags outer "$junos-stacked-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags inner "$junos-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate strict
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile PROF_AUTOSENSE_IPDEMUX
network 10.42.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile PROF_AUTOSENSE_IPDEMUX
network 10.43.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
delimiter "@"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include user-
prefix vlan
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
interface-name
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges session-timeout 600
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-
unit" family inet mac-validate strict

```

```

set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-unit" family inet unnumbered-address lo0.313
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-interface-unit" family inet6 unnumbered-address lo0.313
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht accept any
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht ranges any,any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht accept any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht ranges any
set routing-instances EVPN-VPWS-BNG-2 instance-type evpn-vpws
set routing-instances EVPN-VPWS-BNG-2 protocols evpn interface ps0.0 vpws-service-id local 9999
set routing-instances EVPN-VPWS-BNG-2 protocols evpn interface ps0.0 vpws-service-id remote 1111
set routing-instances EVPN-VPWS-BNG-2 interface ps0.0
set routing-instances EVPN-VPWS-BNG-2 route-distinguisher 172.31.100.4:11
set routing-instances EVPN-VPWS-BNG-2 vrf-target target:65000:11

```

Device ACX1

```

set system host-name ACX1
set system services ssh root-login allow
set chassis network-services enhanced-ip
set system processes dhcp-service traceoptions file dhcp_logfile
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-ADD-DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170

```



```

set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.12.1/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set chassis aggregated-devices ethernet device-count 10
set interfaces ge-0/0/3 gigether-options 802.3ad ae10
set interfaces ae10 flexible-vlan-tagging
set interfaces ae10 encapsulation flexible-ethernet-services
set interfaces ae10 esi 00:11:11:11:11:11:11:11
set interfaces ae10 esi all-active
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 00:00:00:00:00:10
set interfaces ae10 unit 0 encapsulation vlan-ccc
set interfaces ae10 unit 0 vlan-id-list 301-500
set interfaces ae10 unit 0 input-vlan-map push
set interfaces ae10 unit 0 input-vlan-map vlan-id 100
set interfaces ae10 unit 0 output-vlan-map pop
set interfaces lo0 unit 0 family inet address 172.31.100.11/32
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0011.00
set interfaces lo0 unit 0 family mpls
set protocols isis interface all level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
set routing-options router-id 172.31.100.11
set routing-options autonomous-system 65000
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.11
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.4

```

```

set protocols ldp interface all
set protocols mpls interface all
set routing-instances EVPN-VPWS instance-type evpn-vpws
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id local 1111
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id remote 9999
set routing-instances EVPN-VPWS interface ae10.0
set routing-instances EVPN-VPWS route-distinguisher 172.31.100.11:11
set routing-instances EVPN-VPWS vrf-target target:65000:11

```

Device ACX2

```

set system host-name ACX2
set system services ssh root-login allow
set chassis network-services enhanced-ip
set system processes dhcp-service traceoptions file dhcp_logfile
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-ADD-
DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.21.1/30
set interfaces ge-0/0/2 unit 0 family iso

```

```

set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 gigether-options 802.3ad ae10
set interfaces ae10 flexible-vlan-tagging
set interfaces ae10 encapsulation flexible-ethernet-services
set interfaces ae10 esi 00:11:11:11:11:11:11:11
set interfaces ae10 esi all-active
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 00:00:00:00:00:10
set interfaces ae10 unit 0 encapsulation vlan-ccc
set interfaces ae10 unit 0 vlan-id-list 301-500
set interfaces ae10 unit 0 input-vlan-map push
set interfaces ae10 unit 0 input-vlan-map vlan-id 100
set interfaces ae10 unit 0 output-vlan-map pop
set interfaces lo0 unit 0 family inet address 172.31.100.12/32
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0012.00
set interfaces lo0 unit 0 family mpls
set protocols isis interface all level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
set routing-options router-id 172.31.100.12
set routing-options autonomous-system 65000
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.12
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.4
set protocols ldp interface all
set protocols mpls interface all
set routing-instances EVPN-VPWS instance-type evpn-vpws
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id local 1111
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id remote 9999
set routing-instances EVPN-VPWS interface ae10.0
set routing-instances EVPN-VPWS route-distinguisher 172.31.100.12:11
set routing-instances EVPN-VPWS vrf-target target:65000:11

```

Step-by-Step Procedure

Configuring BNG1

1. Log in to the BNG1 device. Ensure that the device is running Junos Release 22.4R1 or later versions.
2. Configure system settings.

```
set system host-name BNG1
set system configuration-database max-db-size 698343424
set system services ssh root-login allow
set chassis network-services enhanced-ip
```

3. Create a group to define common core interfaces configuration such as MTU, hold-time and damping parameters.

```
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED
*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
```

4. Configure the interfaces towards core devices.

```
set interfaces ge-0/0/0 apply-groups GR-CORE-INTF
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
```

```

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.21.2/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls

```

5. Configure the interface towards the vQFX.

```

set interfaces ge-0/0/3 mtu 9192
set interfaces ge-0/0/3 unit 0 family inet address 192.168.100.171/23

```

6. Configure the loopback interface.

```

set interfaces lo0 unit 0 family inet address 172.31.100.3/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0003.00
set interfaces lo0 unit 0 family mpls

```

7. Configure IS-IS protocol in the core network.

```

set protocols isis interface ge-0/0/0.0
set protocols isis interface ge-0/0/1.0
set protocols isis interface ge-0/0/2.0
set protocols isis interface lo0.0 passive
set protocols isis level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept

```

```
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
```

8. Configure routing options.

```
set routing-options router-id 172.31.100.3
set routing-options autonomous-system 65000
```

9. Configure the BGP protocol between the BNG and access devices

```
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.3
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP cluster 1.1.1.1
set protocols bgp group IBGP neighbor 172.31.100.4
set protocols bgp group IBGP neighbor 172.31.100.11
set protocols bgp group IBGP neighbor 172.31.100.12
```

10. Configure LDP and MPLS for all core interfaces.

```
set protocols ldp deaggregate
set protocols ldp transport-address 172.31.100.3
set protocols ldp interface all
set protocols mpls interface all
```

11. Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

```
set system services subscriber-management traceoptions file submgmt.log
set system services subscriber-management traceoptions file size 30m
set system services subscriber-management traceoptions file files 10
set system services subscriber-management traceoptions flag all
set system services subscriber-management gres-route-flush-delay
set system services subscriber-management enable
```

12. Configure tracing options for the general authentication service.

```
set system processes general-authentication-service traceoptions file authd
set system processes general-authentication-service traceoptions file size 10m
set system processes general-authentication-service traceoptions file files 10
set system processes general-authentication-service traceoptions flag all
```

13. Configure system services, including tracing operations and Routing Engine failover, for the main enhanced subscriber management session management process, smg-service.

```
set system processes smg-service traceoptions file smgd
set system processes smg-service traceoptions file size 10m
set system processes smg-service traceoptions file files 10
set system processes smg-service traceoptions level all
set system processes smg-service traceoptions flag all
```

14. Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.

```
set system processes dhcp-service traceoptions file jdhcpd
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions file files 10
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set system processes dhcp-service traceoptions flag all
```

15. Configure the tunnel services and any additional chassis configuration.

```
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis fpc 0 performance-mode
```

16. Configure access profiles for DHCP subscribers.

```
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
```

17. Configure the pseudowire interface to use dynamic stacked VLANs. Also, configure additional interface and VLAN subscription settings. Configure Ethernet Segment Identifier (ESI) for EVPN active-standby multihoming.

```

set chassis pseudowire-service device-count 10
set interfaces ps0 anchor-point lt-0/0/10
set interfaces ps0 flexible-vlan-tagging
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include
delimiter "@"
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include user-
prefix vlan
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include
interface-name
set interfaces ps0 auto-configure stacked-vlan-ranges access-profile no-auth
set interfaces ps0 auto-configure remove-when-no-subscribers
set interfaces ps0 mtu 2022
set interfaces ps0 esi 00:10:00:00:01:00:00:00:10:00
set interfaces ps0 esi single-active
set interfaces ps0 esi df-election-type preference value 1000
set interfaces ps0 unit 0 encapsulation ethernet-ccc

```

18. Configure internet VRF for internet routes.

```

set interfaces lo0 unit 20 description "VRF:internet Loopback"
set interfaces lo0 unit 20 family inet address 172.16.100.3/32 primary
set routing-instances internet instance-type vrf
set routing-instances internet routing-options auto-export
set routing-instances internet interface lo0.20
set routing-instances internet route-distinguisher 172.31.100.3:12
set routing-instances internet vrf-import internet-vrf-import-pol
set routing-instances internet vrf-export internet-vrf-export-pol
set routing-instances internet vrf-table-label
set policy-options policy-statement internet-vrf-export-pol term all then community add
65000:999
set policy-options policy-statement internet-vrf-export-pol term all then accept
set policy-options policy-statement internet-vrf-import-pol term default from community
65000:999
set policy-options policy-statement internet-vrf-import-pol term default then accept
set policy-options policy-statement internet-vrf-import-pol term subs from community
65000:1131
set policy-options policy-statement internet-vrf-import-pol term subs then accept

```



```

set policy-options policy-statement internet-vrf-import-pol term other from community
65000:113
set policy-options policy-statement internet-vrf-import-pol term other then accept
set policy-options community 65000:113 members target:65000:113
set policy-options community 65000:1131 members target:65000:1131
set policy-options community 65000:999 members target:65000:999

```

19. Configure the DHCP local server options on a routing instance. You will configure and apply routing policies for the DHCP subscriber routing instance, create domain maps, firewall filters and dynamic profiles for DHCP subscribers.

```

set interfaces lo0 unit 313 description "VRF:dhcp-sub Loopback"
set interfaces lo0 unit 313 family inet address 172.16.16.3/32
set interfaces lo0 unit 313 family inet address 10.42.0.1/32 primary
set interfaces lo0 unit 313 family inet6 address 2015:cafe:2000::1/128
set routing-instances dhcp-sub instance-type vrf
set routing-instances dhcp-sub routing-options rib dhcp-sub.inet6.0 static route ::/0
next-table internet.inet6.0
set routing-instances dhcp-sub routing-options rib dhcp-sub.inet6.0 static route ::/0 no-
readvertise
set routing-instances dhcp-sub routing-options router-id 172.16.16.3
set routing-instances dhcp-sub routing-options flow term-order standard
set routing-instances dhcp-sub routing-options static route 0.0.0.0/0 next-table
internet.inet.0
set routing-instances dhcp-sub routing-options static route 0.0.0.0/0 no-readvertise
set routing-instances dhcp-sub routing-options static route 10.42.0.0/16 discard
set routing-instances dhcp-sub routing-options static route 10.42.0.0/16 tag 200
set routing-instances dhcp-sub routing-options auto-export
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection failure-action clear-binding-if-interface-up
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd version automatic
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd minimum-interval 30000
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd multiplier 3
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls overrides
client-discover-match incoming-interface
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls overrides
dual-stack dhcp-ds
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls interface
demux0.0

```

```

set routing-instances dhcp-subs system services dhcp-local-server group dhcp-ls interface
ps0.0
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
authentication username-include mac-address
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
on-demand-address-allocation
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
classification-key mac-address
set routing-instances dhcp-subs system services dhcp-local-server dual-stack-group dhcp-ds
protocol-master inet
set routing-instances dhcp-subs system services dhcp-local-server no-stale-timer-refresh
set routing-instances dhcp-subs system services dhcp-local-server stale-timer 60
set routing-instances dhcp-subs access address-assignment high-utilization 80
set routing-instances dhcp-subs access address-assignment abated-utilization 70
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet network 10.42.0.0/16
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet range range1 low 10.42.0.2
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet range range1 high 10.42.255.254
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes maximum-lease-time 600
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes server-identifier 10.42.0.1
set routing-instances dhcp-subs access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes router 10.42.0.1
set routing-instances dhcp-subs access-profile no-auth
set routing-instances dhcp-subs interface lo0.313
set routing-instances dhcp-subs route-distinguisher 172.31.100.3:13
set routing-instances dhcp-subs vrf-import dhcp-subs-vrf-import-pol
set routing-instances dhcp-subs vrf-export dhcp-subs-vrf-export-pol
set routing-instances dhcp-subs vrf-table-label
set policy-options policy-statement dhcp-subs-vrf-export-pol term loop from protocol direct
set policy-options policy-statement dhcp-subs-vrf-export-pol term loop from route-filter
172.16.16.3/32 exact
set policy-options policy-statement dhcp-subs-vrf-export-pol term loop then community add
65000:113
set policy-options policy-statement dhcp-subs-vrf-export-pol term pools from protocol static
set policy-options policy-statement dhcp-subs-vrf-export-pol term pools from tag 200
set policy-options policy-statement dhcp-subs-vrf-export-pol term pools then community add
65000:113
set policy-options policy-statement dhcp-subs-vrf-export-pol term subs then community add
65000:1131

```

```

set policy-options policy-statement dhcp-sub-vrf-export-pol term subs then accept
set policy-options policy-statement dhcp-sub-vrf-import-pol term all from community
65000:111
set policy-options policy-statement dhcp-sub-vrf-import-pol term all then accept
set policy-options community 65000:111 members target:65000:111
set policy-options community 65000:111 members target:65000:11
set access domain map none access-profile no-auth
set access domain map none target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-address
255.255.255.255/32
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-port dhcp
set firewall family inet filter rpf-pass-dhcp term allow-dhcp then accept
set firewall family inet filter rpf-pass-dhcp term default then discard
set system dynamic-profile-options versioning
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" routing-
options access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" actual-
transit-statistics
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" demux-
options underlying-interface "$junos-underlying-interface"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet rpf-check fail-filter rpf-pass-dhcp
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet demux-source $junos-subscriber-ip-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet unnumbered-address lo0.313
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet6 demux-source $junos-subscriber-ipv6-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet6 unnumbered-address lo0.313
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds
dynamic-profile prod-dhcp-base

```

20. Configure and apply dynamic profiles for pseudowire interface.

```

set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix next-hop "$junos-framed-route-nexthop"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles auto-pwht interfaces interface-set "$junos-phy-ifd-interface-set-name"
interface "$junos-interface-ifd-name" unit "$junos-interface-unit"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp restricted
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-id "$junos-vlan-id"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate loose
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags outer "$junos-stacked-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags inner "$junos-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate strict
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile
PROF_AUTOSENSE_IPDEMUX network 10.42.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile
PROF_AUTOSENSE_IPDEMUX network 10.43.0.0/16

```

```

set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
delimiter "@"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
user-prefix vlan
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
interface-name
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges session-timeout 600
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet mac-validate strict
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet unnumbered-address lo0.313
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet6 unnumbered-address lo0.313
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht
accept any
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht
ranges any,any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht accept any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht ranges any

```

21. Configure a routing instance of type evpn-vpws, defining route distinguisher and the VRF target.

```

set routing-instances EVPN-VPWS-BNG-1 instance-type evpn-vpws
set routing-instances EVPN-VPWS-BNG-1 protocols evpn interface ps0.0 vpws-service-id local
9999
set routing-instances EVPN-VPWS-BNG-1 protocols evpn interface ps0.0 vpws-service-id remote
1111
set routing-instances EVPN-VPWS-BNG-1 interface ps0.0
set routing-instances EVPN-VPWS-BNG-1 route-distinguisher 172.31.100.3:11
set routing-instances EVPN-VPWS-BNG-1 vrf-target target:65000:11

```

Configuring BNG2

1. Log in to the BNG2 device. Ensure that the device is running Junos Release 22.4R1 or later versions.

2. Configure system services.

```
set system host-name BNG2
set system configuration-database max-db-size 698343424
set system services ssh root-login allow
set chassis network-services enhanced-ip
```

3. Create a Group to define common core interfaces configuration such as MTU, hold-time and damping parameters.

```
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-
ADD-DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
```

4. Configure the interfaces towards core devices.

```
set interfaces ge-0/0/0 apply-groups GR-CORE-INTF
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.2/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.12.2/30
```

```
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
```

5. Configure the interface towards the vQFX.

```
set interfaces ge-0/0/3 mtu 9192
set interfaces ge-0/0/3 unit 0 family inet address 192.168.100.171/23
```

6. Configure the loopback interface for use in the subscriber management access network.

```
set interfaces lo0 unit 0 family inet address 172.31.100.4/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0004.00
set interfaces lo0 unit 0 family mpls
```

7. Configure IS-IS protocol in the core network.

```
set protocols isis interface ge-0/0/0.0
set protocols isis interface ge-0/0/1.0
set protocols isis interface ge-0/0/2.0
set protocols isis interface lo0.0 passive
set protocols isis level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT
```

8. Configure routing options.

```
set routing-options router-id 172.31.100.4
set routing-options autonomous-system 65000
```

9. Configure the BGP protocol between the BNG and access devices.

```
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.4
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP cluster 2.2.2.2
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.11
set protocols bgp group IBGP neighbor 172.31.100.12
```

10. Configure LDP and MPLS for all core interfaces.

```
set protocols ldp deaggregate
set protocols ldp transport-address 172.31.100.4
set protocols ldp interface all
set protocols mpls interface all
```

11. Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

```
set system services subscriber-management traceoptions file submgmt.log
set system services subscriber-management traceoptions file size 30m
set system services subscriber-management traceoptions file files 10
set system services subscriber-management traceoptions flag all
set system services subscriber-management gres-route-flush-delay
set system services subscriber-management enable
```

12. Configure tracing options for the general authentication service.

```
set system processes general-authentication-service traceoptions file authd
set system processes general-authentication-service traceoptions file size 10m
set system processes general-authentication-service traceoptions file files 10
set system processes general-authentication-service traceoptions flag all
```


13. Configure system services, including tracing operations and Routing Engine failover, for the main enhanced subscriber management session management process, smg-service.

```
set system processes smg-service traceoptions file smgd
set system processes smg-service traceoptions file size 10m
set system processes smg-service traceoptions file files 10
set system processes smg-service traceoptions level all
set system processes smg-service traceoptions flag all
```

14. Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.

```
set system processes dhcp-service traceoptions file jdhcpd
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions file files 10
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
set system processes dhcp-service traceoptions flag all
```

15. Configure the tunnel services and any additional chassis configuration.

```
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis fpc 0 performance-mode
```

16. Configure access profiles for DHCP subscribers.

```
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
```

17. Configure the pseudowire interface to use dynamic stacked VLANs. Also, configure additional interface and VLAN subscription settings. Configure Ethernet Segment Identifier (ESI) for EVPN active-standby multihoming.

```
set chassis pseudowire-service device-count 10
set interfaces ps0 anchor-point lt-0/0/10
set interfaces ps0 flexible-vlan-tagging
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include
```

```

delimiter "@"
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include user-
prefix vlan
set interfaces ps0 auto-configure stacked-vlan-ranges authentication username-include
interface-name
set interfaces ps0 auto-configure stacked-vlan-ranges access-profile no-auth
set interfaces ps0 auto-configure remove-when-no-subscribers
set interfaces ps0 mtu 2022
set interfaces ps0 esi 00:10:00:00:01:00:00:00:10:00
set interfaces ps0 esi single-active
set interfaces ps0 esi df-election-type preference value 999
set interfaces ps0 unit 0 encapsulation ethernet-ccc

```

18. Configure internet VRF for internet routes.

```

set interfaces lo0 unit 20 description "VRF:internet Loopback"
set interfaces lo0 unit 20 family inet address 172.16.100.4/32 primary
set routing-instances internet instance-type vrf
set routing-instances internet routing-options auto-export
set routing-instances internet interface lo0.20
set routing-instances internet route-distinguisher 172.31.100.4:12
set routing-instances internet vrf-import internet-vrf-import-pol
set routing-instances internet vrf-export internet-vrf-export-pol
set routing-instances internet vrf-table-label
set policy-options policy-statement internet-vrf-export-pol term all then community add
65000:999
set policy-options policy-statement internet-vrf-export-pol term all then accept
set policy-options policy-statement internet-vrf-import-pol term default from community
65000:999
set policy-options policy-statement internet-vrf-import-pol term default then accept
set policy-options policy-statement internet-vrf-import-pol term subs from community
65000:1131
set policy-options policy-statement internet-vrf-import-pol term subs then accept
set policy-options policy-statement internet-vrf-import-pol term other from community
65000:113
set policy-options policy-statement internet-vrf-import-pol term other then accept
set policy-options community 65000:113 members target:65000:113
set policy-options community 65000:1131 members target:65000:1131
set policy-options community 65000:999 members target:65000:999

```

19. Configure the DHCP local server options on a routing instance. You will configure and apply routing policies for the DHCP subscriber routing instance, create domain maps, firewall filters and dynamic profiles for DHCP subscribers.

```

set interfaces lo0 unit 313 description "VRF:dhcp-sub Loopback"
set interfaces lo0 unit 313 family inet address 172.16.16.4/32
set interfaces lo0 unit 313 family inet address 10.43.0.1/32 primary
set interfaces lo0 unit 313 family inet6 address 2016:cafe:2000::1/128
set routing-instances dhcp-sub instance-type vrf
set routing-instances dhcp-sub routing-options rib dhcp-sub.inet6.0 static route ::/0
next-table internet.inet6.0
set routing-instances dhcp-sub routing-options rib dhcp-sub.inet6.0 static route ::/0 no-
readvertise
set routing-instances dhcp-sub routing-options router-id 172.16.16.4
set routing-instances dhcp-sub routing-options flow term-order standard
set routing-instances dhcp-sub routing-options static route 0.0.0.0/0 next-table
internet.inet.0
set routing-instances dhcp-sub routing-options static route 0.0.0.0/0 no-readvertise
set routing-instances dhcp-sub routing-options static route 10.43.0.0/16 discard
set routing-instances dhcp-sub routing-options static route 10.43.0.0/16 tag 200
set routing-instances dhcp-sub routing-options auto-export
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection failure-action clear-binding-if-interface-up
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd version automatic
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd minimum-interval 30000
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls liveness-
detection method bfd multiplier 3
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls overrides
client-discover-match incoming-interface
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls overrides
dual-stack dhcp-ds
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls interface
demux0.0
set routing-instances dhcp-sub system services dhcp-local-server group dhcp-ls interface
ps0.0
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds
authentication username-include mac-address
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds
on-demand-address-allocation
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds

```

```

classification-key mac-address
set routing-instances dhcp-subsystem services dhcp-local-server dual-stack-group dhcp-ds
protocol-master inet
set routing-instances dhcp-subsystem services dhcp-local-server no-stale-timer-refresh
set routing-instances dhcp-subsystem services dhcp-local-server stale-timer 60
set routing-instances dhcp-subsystem access address-assignment high-utilization 80
set routing-instances dhcp-subsystem access address-assignment abated-utilization 70
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet network 10.43.0.0/16
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet range range1 low 10.43.0.2
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet range range1 high 10.43.254.254
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes maximum-lease-time 600
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes server-identifier 10.43.0.1
set routing-instances dhcp-subsystem access address-assignment pool ttt-fttx-res-ipv4-pool-0
family inet dhcp-attributes router 10.43.0.1
set routing-instances dhcp-subsystem access-profile no-auth
set routing-instances dhcp-subsystem interface lo0.313
set routing-instances dhcp-subsystem route-distinguisher 172.31.100.4:13
set routing-instances dhcp-subsystem vrf-import dhcp-subsystem-vrf-import-pol
set routing-instances dhcp-subsystem vrf-export dhcp-subsystem-vrf-export-pol
set routing-instances dhcp-subsystem vrf-table-label
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop from protocol direct
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop from route-filter
172.16.16.4/32 exact
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term loop then community add
65000:113
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools from protocol static
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools from tag 200
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term pools then community add
65000:113
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term subs then community add
65000:1131
set policy-options policy-statement dhcp-subsystem-vrf-export-pol term subs then accept
set policy-options policy-statement dhcp-subsystem-vrf-import-pol term all from community
65000:111
set policy-options policy-statement dhcp-subsystem-vrf-import-pol term all then accept
set policy-options community 65000:111 members target:65000:111
set policy-options community 65000:111 members target:65000:11
set policy-options community 65000:113 members target:65000:113

```

```

set policy-options community 65000:1131 members target:65000:1131
set access-profile no-auth
set access profile no-auth authentication-order none
set access address-protection
set access domain map none access-profile no-auth
set access domain map none target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set access domain map ps0.* access-profile no-auth
set access domain map ps0.* target-routing-instance dhcp-sub
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-address
255.255.255.255/32
set firewall family inet filter rpf-pass-dhcp term allow-dhcp from destination-port dhcp
set firewall family inet filter rpf-pass-dhcp term allow-dhcp then accept
set firewall family inet filter rpf-pass-dhcp term default then discard
set system dynamic-profile-options versioning
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles prod-dhcp-base routing-instances "$junos-routing-instance" routing-
options access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" actual-
transit-statistics
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" demux-
options underlying-interface "$junos-underlying-interface"
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet rpf-check fail-filter rpf-pass-dhcp
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet demux-source $junos-subscriber-ip-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet unnumbered-address lo0.313
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet6 demux-source $junos-subscriber-ipv6-address
set dynamic-profiles prod-dhcp-base interfaces demux0 unit "$junos-interface-unit" family
inet6 unnumbered-address lo0.313
set routing-instances dhcp-sub system services dhcp-local-server dual-stack-group dhcp-ds
dynamic-profile prod-dhcp-base

```

20. Configure and apply dynamic profiles for pseudowire interface.

```

set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" interface
"$junos-interface-name"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options

```

```

access route $junos-framed-route-ip-address-prefix next-hop "$junos-framed-route-nexthop"
set dynamic-profiles auto-pwht routing-instances "$junos-routing-instance" routing-options
access route $junos-framed-route-ip-address-prefix metric "$junos-framed-route-cost"
set dynamic-profiles auto-pwht interfaces interface-set "$junos-phy-ifd-interface-set-name"
interface "$junos-interface-ifd-name" unit "$junos-interface-unit"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp restricted
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-id "$junos-vlan-id"
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate loose
set dynamic-profiles auto-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" no-traps
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" proxy-arp
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags outer "$junos-stacked-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" vlan-tags inner "$junos-vlan-id"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet mac-validate strict
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet unnumbered-address "$junos-loopback-interface"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile
PROF_AUTOSENSE_IPDEMUX network 10.42.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges dynamic-profile
PROF_AUTOSENSE_IPDEMUX network 10.43.0.0/16
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
delimiter "@"
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
user-prefix vlan

```

```

set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges authentication username-include
interface-name
set dynamic-profiles auto-stacked-pwht interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" family inet auto-configure address-ranges session-timeout 600
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX routing-instances "$junos-routing-instance"
interface "$junos-interface-name"
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet mac-validate strict
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet unnumbered-address lo0.313
set dynamic-profiles PROF_AUTOSENSE_IPDEMUX interfaces demux0 unit "$junos-underlying-
interface-unit" family inet6 unnumbered-address lo0.313
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht
accept any
set interfaces ps0 auto-configure stacked-vlan-ranges dynamic-profile auto-stacked-pwht
ranges any,any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht accept any
set interfaces ps0 auto-configure vlan-ranges dynamic-profile auto-pwht ranges any

```

21. Configure a routing instance of type evpn-vpws, defining route distinguisher and the VRF target.

```

set routing-instances EVPN-VPWS-BNG-2 instance-type evpn-vpws
set routing-instances EVPN-VPWS-BNG-2 protocols evpn interface ps0.0 vpws-service-id local
9999
set routing-instances EVPN-VPWS-BNG-2 protocols evpn interface ps0.0 vpws-service-id remote
1111
set routing-instances EVPN-VPWS-BNG-2 interface ps0.0
set routing-instances EVPN-VPWS-BNG-2 route-distinguisher 172.31.100.4:11
set routing-instances EVPN-VPWS-BNG-2 vrf-target target:65000:11

```

Configuring ACX1

1. Log in to the ACX1 device.
2. Configure System settings and DHCP service settings.

```

set system host-name ACX1
set system services ssh root-login allow
set chassis network-services enhanced-ip
set system processes dhcp-service traceoptions file dhcp_logfile

```

```
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
```

3. Create a Group to define common core interfaces configuration such as MTU, hold-time and damping parameters.

```
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-
ADD-DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
```

4. Configure the interfaces towards the core devices

```
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.12.1/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
```

5. Configure aggregate interface with appropriate VLAN and Ethernet Segment Identifier (ESI) configuration.

```
set chassis aggregated-devices ethernet device-count 10
set interfaces ge-0/0/3 together-options 802.3ad ae10
```



```

set interfaces ae10 flexible-vlan-tagging
set interfaces ae10 encapsulation flexible-ethernet-services
set interfaces ae10 esi 00:11:11:11:11:11:11:11:11
set interfaces ae10 esi all-active
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 00:00:00:00:00:10
set interfaces ae10 unit 0 encapsulation vlan-ccc
set interfaces ae10 unit 0 vlan-id-list 301-500
set interfaces ae10 unit 0 input-vlan-map push
set interfaces ae10 unit 0 input-vlan-map vlan-id 100
set interfaces ae10 unit 0 output-vlan-map pop

```

6. Configure the loopback interfaces.

```

set interfaces lo0 unit 0 family inet address 172.31.100.11/32
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0011.00
set interfaces lo0 unit 0 family mpls

```

7. Configure the IS-IS protocol for the core network.

```

set protocols isis interface all level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT

```

8. Configure routing options.

```

set routing-options router-id 172.31.100.11
set routing-options autonomous-system 65000

```

9. Configure the BGP protocol between the access devices and BNG.

```
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.11
set protocols bgp group IBGP family evpn signaling
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.4
```

10. Configure the LDP and MPLS protocols on the core interfaces.

```
set protocols ldp interface all
set protocols mpls interface all
```

11. Configure a routing instance of type evpn-vpws, defining route distinguisher and the VRF target.

```
set routing-instances EVPN-VPWS instance-type evpn-vpws
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id local 1111
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id remote 9999
set routing-instances EVPN-VPWS interface ae10.0
set routing-instances EVPN-VPWS route-distinguisher 172.31.100.11:11
set routing-instances EVPN-VPWS vrf-target target:65000:11
```

Configuring ACX2

1. Log in to the ACX2 device.
2. Configure System Settings and DHCP service settings.

```
set system host-name ACX2
set system services ssh root-login allow
set chassis network-services enhanced-ip
set system processes dhcp-service traceoptions file dhcp_logfile
set system processes dhcp-service traceoptions file size 10m
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag packet
```

3. Create a Group to define common core interfaces configuration such as MTU, hold-time and damping parameters.

```
set groups GR-CORE-INTF interfaces <*> description *****GR-CORE-INTF-SETTINGS-APPLIED-
ADD-DESCRIPTION*****
set groups GR-CORE-INTF interfaces <*> traps
set groups GR-CORE-INTF interfaces <*> mtu 9192
set groups GR-CORE-INTF interfaces <*> hold-time up 2000
set groups GR-CORE-INTF interfaces <*> hold-time down 0
set groups GR-CORE-INTF interfaces <*> damping half-life 30
set groups GR-CORE-INTF interfaces <*> damping max-suppress 600
set groups GR-CORE-INTF interfaces <*> damping reuse 250
set groups GR-CORE-INTF interfaces <*> damping suppress 2000
set groups GR-CORE-INTF interfaces <*> damping enable
set groups GR-CORE-INTF interfaces <*> unit 0 traps
set groups GR-CORE-INTF interfaces <*> unit 0 family inet mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family iso mtu 9106
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls mtu 9170
set groups GR-CORE-INTF interfaces <*> unit 0 family mpls maximum-labels 5
```

4. Configure the interfaces towards the core devices.

```
set interfaces ge-0/0/1 apply-groups GR-CORE-INTF
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 apply-groups GR-CORE-INTF
set interfaces ge-0/0/2 unit 0 family inet address 10.1.21.1/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
```

5. Configure aggregate interface with appropriate VLAN and Ethernet Segment Identifier (ESI) configuration.

```
set interfaces ge-0/0/3 gigether-options 802.3ad ae10
set interfaces ae10 flexible-vlan-tagging
set interfaces ae10 encapsulation flexible-ethernet-services
set interfaces ae10 esi 00:11:11:11:11:11:11:11:11
set interfaces ae10 esi all-active
set interfaces ae10 aggregated-ether-options lacp active
```

```

set interfaces ae10 aggregated-ether-options lacp system-id 00:00:00:00:00:10
set interfaces ae10 unit 0 encapsulation vlan-ccc
set interfaces ae10 unit 0 vlan-id-list 301-500
set interfaces ae10 unit 0 input-vlan-map push
set interfaces ae10 unit 0 input-vlan-map vlan-id 100
set interfaces ae10 unit 0 output-vlan-map pop

```

6. Configure the loopback interfaces.

```

set interfaces lo0 unit 0 family inet address 172.31.100.12/32
set interfaces lo0 unit 0 family iso address 49.0001.1000.0000.0012.00
set interfaces lo0 unit 0 family mpls

```

7. Configure the IS-IS protocol for the core network.

```

set protocols isis interface all level 1 disable
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from interface lo0.0
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK from route-filter
172.31.0.0/16 prefix-length-range /32-/32
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then tag 101
set policy-options policy-statement PL-ISIS-EXPORT term LOCAL-LOOPBACK then accept
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES from protocol direct
set policy-options policy-statement PL-ISIS-EXPORT term DIRECT-ROUTES then accept
set policy-options policy-statement PL-ISIS-EXPORT then reject
set protocols isis export PL-ISIS-EXPORT

```

8. Configure routing options.

```

set routing-options router-id 172.31.100.12
set routing-options autonomous-system 65000

```

9. Configure the BGP protocol between the access devices and BNG.

```

set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 172.31.100.12
set protocols bgp group IBGP family evpn signaling

```

```
set protocols bgp group IBGP neighbor 172.31.100.3
set protocols bgp group IBGP neighbor 172.31.100.4
```

10. Configure the LDP and MPLS protocols on the core interfaces.

```
set protocols ldp interface all
set protocols mpls interface all
```

11. Configure a routing instance of type evpn-vpws, defining route distinguisher and the VRF target.

```
set routing-instances EVPN-VPWS instance-type evpn-vpws
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id local 1111
set routing-instances EVPN-VPWS protocols evpn interface ae10.0 vpws-service-id remote 9999
set routing-instances EVPN-VPWS interface ae10.0
set routing-instances EVPN-VPWS route-distinguisher 172.31.100.12:11
set routing-instances EVPN-VPWS vrf-target target:65000:11
```

Verification

IN THIS SECTION

- Purpose | **1614**
- Action | **1615**
- Meaning | **1619**

Purpose

You can verify the configuration by issuing the `show evpn vpws-instance` command on the BNG and access devices to view the details of the VPWS instance of the EVPN.

Action

- Verifying on the BNG1 device.

```
show evpn vpws-instance
Instance: EVPN-VPWS-BNG-1, Instance type: EVPN VPWS, Encapsulation type: MPLS
Route Distinguisher: 172.31.100.3:11
Number of local interfaces: 1 (1 up)

Interface name  ESI                               Mode      Role      Status
Control-Word   Flow-Label-Tx   Flow-Label-Rx
ps0.0          00:10:00:00:01:00:00:00:10:00 single-active Primary   Up
No             No              No
Local SID: 9999 Advertised Label: 46
PE addr      ESI                               Label  End.Dx2 SID
Mode         Role    TS                               Status
172.31.100.4 00:10:00:00:01:00:00:00:10:00 64          single-
active Backup 2023-07-24 11:08:27.958 Resolved
Remote SID: 1111
PE addr      ESI                               Label  End.Dx2 SID
Mode         Role    TS                               Status
172.31.100.11 00:11:11:11:11:11:11:11:11:11 300112      all-
active Primary 2023-07-24 11:08:17.689 Resolved
172.31.100.12 00:11:11:11:11:11:11:11:11:11 300096      all-
active Primary 2023-07-24 11:08:17.689 Resolved
Number of protect interfaces: 0

Fast Convergence Information
ESI: 00:10:00:00:01:00:00:00:10:00 Number of PE nodes: 1
PE: 172.31.100.4
Advertised SID: 9999

Fast Convergence Information
ESI: 00:11:11:11:11:11:11:11:11:11 Number of PE nodes: 2
PE: 172.31.100.11
Advertised SID: 1111
PE: 172.31.100.12
Advertised SID: 1111

DF Election Information for Single-Active ESI
ESI: 00:10:00:00:01:00:00:00:10:00
DF Election Algorithm: Preference based
```

Primary PE: 172.31.100.3, Preference: 1000
 Backup PE: 172.31.100.4, Preference: 999
 Last DF Election: 2023-07-24 11:08:27

- Verifying on the BNG2 device.

```
show evpn vpws-instance
```

Instance: EVPN-VPWS-BNG-2, Instance type: EVPN VPWS, Encapsulation type: MPLS

Route Distinguisher: 172.31.100.4:11

Number of local interfaces: 1 (1 up)

Interface name	ESI	Mode	Role	Status
Control-Word	Flow-Label-Tx	Flow-Label-Rx		
ps0.0	00:10:00:00:01:00:00:00:10:00	single-active	Backup	Up
No	No	No		
Local SID: 9999 Advertised Label: 64				
PE addr	ESI	Label	End.Dx2	SID
Mode	Role	TS	Status	
172.31.100.3	00:10:00:00:01:00:00:00:10:00	46		single-
active	Primary	2023-07-24 11:08:27.750	Resolved	
Remote SID: 1111				
PE addr	ESI	Label	End.Dx2	SID
Mode	Role	TS	Status	
172.31.100.11	00:11:11:11:11:11:11:11:11:11	300112		all-
active	Primary	2023-07-24 10:59:26.019	Resolved	
172.31.100.12	00:11:11:11:11:11:11:11:11:11	300096		all-
active	Primary	2023-07-24 10:59:26.317	Resolved	
Number of protect interfaces: 0				

Fast Convergence Information

ESI: 00:10:00:00:01:00:00:00:10:00 Number of PE nodes: 1

PE: 172.31.100.3

Advertised SID: 9999

Fast Convergence Information

ESI: 00:11:11:11:11:11:11:11:11:11 Number of PE nodes: 2

PE: 172.31.100.11

Advertised SID: 1111

PE: 172.31.100.12

Advertised SID: 1111

DF Election Information for Single-Active ESI

```

ESI: 00:10:00:00:01:00:00:00:10:00
DF Election Algorithm: Preference based
Primary PE: 172.31.100.3, Preference: 1000
Backup PE: 172.31.100.4, Preference: 999
Last DF Election: 2023-07-24 11:08:27

```

- Verifying on ACX1 device.

```

show evpn vpws-instance
Instance: EVPN-VPWS, Instance type: EVPN VPWS, Encapsulation type: MPLS
Route Distinguisher: 172.31.100.11:11
Number of local interfaces: 1 (1 up)

Interface name  ESI                               Mode      Role      Status
Control-Word   Flow-Label-Tx   Flow-Label-Rx
ae10.0         00:11:11:11:11:11:11:11:11:11:11 all-active Primary   Up
No             No              No
Local SID: 1111 Advertised Label: 300112
PE addr      ESI                               Label  End.Dx2  SID
Mode         Role    TS                               Status
172.31.100.12 00:11:11:11:11:11:11:11:11:11:11 300096          all-
active      Primary 2023-07-24 10:59:26.522 Resolved
Remote SID: 9999
PE addr      ESI                               Label  End.Dx2  SID
Mode         Role    TS                               Status
172.31.100.4  00:10:00:00:01:00:00:00:10:00 64          single-
active      Backup  2023-07-24 11:08:27.963 Resolved
172.31.100.3  00:10:00:00:01:00:00:00:10:00 46          single-
active      Primary 2023-07-24 11:08:27.749 Resolved
Number of protect interfaces: 0

Fast Convergence Information
ESI: 00:10:00:00:01:00:00:00:10:00 Number of PE nodes: 2
PE: 172.31.100.4
Advertised SID: 9999
PE: 172.31.100.3
Advertised SID: 9999

Fast Convergence Information
ESI: 00:11:11:11:11:11:11:11:11:11:11 Number of PE nodes: 1

```



```
PE: 172.31.100.12
Advertised SID: 1111
```

- Verifying on the ACX2 device.

```
show evpn vpws-instance
Instance: EVPN-VPWS, Instance type: EVPN VPWS, Encapsulation type: MPLS
Route Distinguisher: 172.31.100.12:11
Number of local interfaces: 1 (1 up)
```

Interface name	ESI	Mode	Role	Status
Control-Word	Flow-Label-Tx	Flow-Label-Rx		
ae10.0	00:11:11:11:11:11:11:11:11	all-active	Primary	Up
No	No	No		
Local SID: 1111 Advertised Label: 300096				
PE addr	ESI	Label	End.Dx2	SID
Mode	Role	TS	Status	
172.31.100.11	00:11:11:11:11:11:11:11:11	300112		all-
active	Primary	2023-07-24 10:59:26.224	Resolved	
Remote SID: 9999				
PE addr	ESI	Label	End.Dx2	SID
Mode	Role	TS	Status	
172.31.100.4	00:10:00:00:01:00:00:00:10:00	64		single-
active	Backup	2023-07-24 11:08:27.960	Resolved	
172.31.100.3	00:10:00:00:01:00:00:00:10:00	46		single-
active	Primary	2023-07-24 11:08:27.750	Resolved	
Number of protect interfaces: 0				
Fast Convergence Information				
ESI: 00:10:00:00:01:00:00:00:10:00 Number of PE nodes: 2				
PE: 172.31.100.4				
Advertised SID: 9999				
PE: 172.31.100.3				
Advertised SID: 9999				
Fast Convergence Information				
ESI: 00:11:11:11:11:11:11:11:11 Number of PE nodes: 1				
PE: 172.31.100.11				
Advertised SID: 1111				

Meaning

The BNG devices are configured in an active-standby multihoming mode. In a steady state, all flows are directed towards BNG1, which is the primary device between BNG1 and BNG2. In case BNG1 encounters a failure, BNG2 becomes the primary designated forwarder. The access devices are configured in active-active multihoming, load balancing all the traffic from the CE devices.

RELATED DOCUMENTATION

auto-configure (IPv4)

auto-configure (IPv6)

9

PART

Access Node Control Protocol and the ANCP Agent for Subscriber Services

- Access Node Control Protocol and the ANCP Agent for Subscriber Services | **1621**
-

CHAPTER 12

Access Node Control Protocol and the ANCP Agent for Subscriber Services

IN THIS CHAPTER

- [ANCP Agent Neighbors and Operations | 1621](#)
- [ANCP Agent Traffic Shaping and CoS | 1681](#)
- [ANCP Agent and AAA | 1700](#)
- [ANCP Monitoring and Management | 1715](#)
- [Tracing ANCP Events for Troubleshooting | 1722](#)

ANCP Agent Neighbors and Operations

IN THIS SECTION

- [ANCP and the ANCP Agent Overview | 1622](#)
- [ANCP Operations in Different Network Configurations | 1632](#)
- [Configuring the ANCP Agent | 1642](#)
- [Configuring ANCP Neighbors | 1644](#)
- [Associating an Access Node with Subscribers for ANCP Agent Operations | 1645](#)
- [Specifying the Interval Between ANCP Adjacency Messages | 1646](#)
- [Specifying the Maximum Number of Discovery Table Entries | 1647](#)
- [Configuring the ANCP Agent for Backward Compatibility | 1647](#)
- [Specifying How Long Processes Wait for the ANCP Agent Restart to Complete | 1648](#)
- [Configuring the ANCP Agent to Learn ANCP Partition IDs | 1649](#)
- [Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet | 1650](#)

ANCP and the ANCP Agent Overview

IN THIS SECTION

- [Overview | 1622](#)
- [Topology Discovery | 1623](#)
- [Subscriber Services | 1623](#)
- [ANCP Interfaces and Access Loop Circuit Identifiers | 1624](#)
- [Mapping Access Lines to Interfaces and Interface Sets | 1625](#)
- [ANCP Neighbors | 1626](#)
- [Partitions | 1628](#)
- [Adjacency Update Messages | 1629](#)
- [Generic Response Messages and Result Codes | 1629](#)
- [Benefits of Access Node Control Protocol | 1631](#)

This topic describes the Access Node Control Protocol (ANCP) and the *ANCP agent*. The ANCP agent is the Junos OS process that manages subscriber access lines with ANCP. The agent monitors subscriber access lines, reports subscriber traffic rates on the access lines between the subscribers and the access nodes, and modifies the traffic rates, all in support of CoS traffic shaping.

Overview

ANCP acts as a control plane between a service-oriented Layer 3 edge device and a Layer 2 access node. The access nodes—ANCP *neighbors*—are network devices that terminate access loops from subscribers; for DSL access loops, the access node is a DSL access multiplexer (DSLAM). Queuing and scheduling mechanisms for subscriber traffic must avoid congestion within the access network while contending with multiple flows and distinct CoS requirements. These mechanisms require the edge device—a router acting as a broadband network gateway (BNG), often also called a network access server (NAS)—to provide information about the access network and subscriber traffic.

The ANCP agent can map an access line to an interface or interface set either statically or dynamically. The agent provides that information to both CoS and AAA. The agent passes on to both CoS and AAA the traffic shaping attributes for each subscriber access line that the access node sent to the ANCP agent. In addition, the agent sends to AAA all DSL Forum attributes that were sent by the access node. AAA can use these attributes during RADIUS accounting and authentication for both DHCP IP demux and PPPoE subscriber sessions. The traffic rates can also be used for shaping L2TP tunnel traffic.

You can monitor ANCP agent events and operations by including the `traceoptions` statement at the `[edit protocols ancp]` hierarchy level.

Junos OS supports Class of Service (CoS) traffic shaping on the following interface types for ANCP:

- Static VLAN interfaces, except those created by Extensible Subscriber Services Manager (ESSM)
- Static VLAN demux interfaces, except those created by ESSM
- Static interface sets, including those created by ESSM
- Dynamic interface sets
- Dynamic VLAN-tagged interface sets
- Dynamic agent circuit identifier (ACI) interface sets, also known as ACI sets or ACI VLANs
- Dynamic PPPoE and DHCP IP demux subscriber interfaces
- Dynamic VLAN demux interfaces with Ethernet-VPLS encapsulation

ANCP was developed as an extension of *RFC 3292, General Switch Management Protocol (GSMP) V3*, but is now defined in *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*.

Topology Discovery

The router uses topology discovery to collect information from the access node. The information includes the following:

- Topology of the access network
- DSL line state
- Actual upstream and downstream net data rates of a synchronized DSL link
- Maximum attainable upstream and downstream net data rates
- Interleaving delay

Subscriber Services

The router receives the service profile for the subscribers from a RADIUS server. Most of the services are enforced by the router itself. The router shapes the aggregate egress traffic to subscribers based on the local loop throughput reported by the DSLAM. This traffic shaping optimizes traffic flow while avoiding traffic drops in the access node.

Some service attributes, such as interleaving delay and multicast channel information, are enforced at the access node. The ANCP agent provides the line configuration mechanism that the edge device can

use to pass the line configuration to the access nodes. Typically, multiple profiles are provisioned on the access node. The router instructs the access node which profile to use for a given subscriber.

Subscribers typically receive some combination of voice, data, and video services. Each service can be provisioned on a VLAN. A subscriber might receive only a single service over a single VLAN configured on a *logical interface*. A group of VLANs carrying services to a subscriber is an *interface set*.

Subscribers have operational states, but they do not have administrative states because they cannot be configured in the CLI.

Subscribers have one of the following operational states which represent the DSL line state as it is reported in the ANCP Port Up and Port Down messages sent by an access node:

- Idle—Ports are not configured and the subscriber cannot log in.
- Silent—Ports are configured and the subscriber is connected, but the DSL modem is not ready to transfer data.
- Showtime—Ports are configured, the subscriber is connected, and the DSL modem is online and ready to transfer data.



NOTE: For information about ANCP for business subscribers and services, see *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*.

ANCP Interfaces and Access Loop Circuit Identifiers

The access loop or access line in an ANCP topology consists of the physical elements between the subscriber device (CPE) and the access node. An identifier associated with the access loop serves to identify the subscriber as well. This identifier is an alphanumeric string that actually identifies the interface on the DSLAM from which subscriber requests originate. It can be referred to by various names.

- In ANCP messages, a TLV carries the access loop circuit ID, also referred to as the access line identifier, access loop circuit identifier, or access identifier.
- DHCP discovery packets can identify the line with the Agent Circuit ID suboption in the Option 82 field.
- PPPoE discovery packets can identify the line with the Agent-Circuit-ID subattribute in the DSL Forum vendor-specific tag.

Each of these identifiers is abbreviated as ACI. When the ANCP agent receives a port management message from an access node, it uses the access loop circuit identifier contained in the message to determine which logical interface or interface set corresponds to the subscriber.

You can associate an identifier with an ANCP access line by static configuration. When you configure a logical interface by specifying the interface name at the [edit protocols ancp interfaces] hierarchy level, include the `access-identifier` statement to associate the access loop circuit identifier with the interface. When you configure an interface set by including the `interface-set` statement at the [edit protocols ancp interfaces] hierarchy level, associate the access loop circuit identifier with the interface set by including the `access-identifier` statement at the [edit protocols ancp interfaces interface-set *interface-set-name*] hierarchy level.

When the DHCP or PPPoE discovery packet includes an ACI, the ANCP agent can dynamically map the ACI to the subscriber interface or interface set. VLANs for the subscribers are created according to a dynamic profile; these are called agent circuit identifier-based or ACI-based dynamic VLANs.

ANCP agent support for RADIUS authentication and accounting requires that both static and dynamic ACIs must be unique across the network. No two interfaces across multiple neighbors (access nodes) can share the same identifier. The DHCP and PPPoE processes do not have information about the access node IP addresses and consequently cannot distinguish between duplicate identifiers. This situation prevents the AAA services framework from correlating a DHCP or PPPoE client session with an access line for RADIUS authentication and accounting.

Mapping Access Lines to Interfaces and Interface Sets

The ANCP agent maps the ACI for subscriber access lines to an interface or interface set to apply DSL attributes received from the access node to CoS traffic shaping for the access line. The access line mapping can be statically configured with the `access-identifier` statement, or dynamically derived during subscriber authentication. Static mapping always supersedes dynamic mapping.

The ANCP agent can remap an access line to a different interface or interface set than its original mapping. Remapping can also be static or dynamic. For example, an access line might be first dynamically mapped to a subscriber interface and then statically configured to an interface set.

You can statically configure mapping with the statement only for interface and interface set types that have configured or deterministic names:

- Static VLAN interfaces
- Static VLAN demux interfaces
- Static interface sets
- Dynamic interface sets
- Dynamic VLAN-tagged interface sets

Static configuration with the statement is required for mapping an access line to static interface sets, dynamic interface sets, and dynamic VLAN-tagged interface sets. This is true regardless of the presence

of an ACI in the PPPoE or DHCP IP demux subscriber's discovery packet, because the use of the ACI is irrelevant to the creation of these types of interface sets.

You cannot statically configure mapping with the statement for the following interface and interface set types, because they have nondeterministic, automatically generated names:

- Dynamic VLAN demux interfaces
- Dynamic ACI interface sets (ACI VLANs)
- Dynamic PPPoE and DHCP IP demux subscriber interfaces

In the context of Layer 2 wholesale services, the ANCP agent can map access lines to dynamic VLAN demux interfaces that have Ethernet-VPLS encapsulation. The ANCP agent triggers the creation of these interfaces with the ANCP Port UP message, which always includes the ACI for the access line. The agent can then dynamically map the interface to an access line for CoS traffic shaping.

Dynamic mapping works as follows:

- If the subscriber interface is a member of an interface set, the ANCP agent maps the ACI for the access line to the interface set.
- If the subscriber interface is not a member of an interface set, the ANCP agent maps the ACI for the access line to the subscriber interface.

The ANCP agent does not support static or dynamic mapping for the following interface types, regardless of the presence of the access line's ACI in the subscriber's discovery packet:

- Static VLAN interfaces created by ESSM.
- Static VLAN demux interfaces created by ESSM.
- Dynamic VLAN interfaces.
- Dynamic VLAN demux interfaces that do not have Ethernet-VPLS encapsulation.

ANCP Neighbors

The ANCP agent can report traffic only for access nodes that are configured as ANCP neighbors (also referred to as ANCP peers). Neighbors can establish TCP connections with the router. Include the neighbor statement at the [edit protocols ancp] hierarchy level to configure an access node as an ANCP neighbor.

The ANCP agent exchanges adjacency messages with neighbors. If an adjacency message is not received from a neighbor within the expected period, then the neighbor is considered to be down and is disconnected. You can adjust how long the ANCP agent waits for adjacency messages from all neighbors by including the adjacency-timer statement at the [edit protocols ancp] hierarchy level. The interval between adjacency messages is negotiated between router and the neighbor during adjacency

establishment. The larger of two timer values—either the value received in the ANCP SYN message or the configured value—is selected. Loss of synchronization between the router and a neighbor is declared when no valid messages are received for a period of time that exceeds three times the negotiated value.



NOTE: The ANCP TCP connection is not established and consequently ANCP neighbors do not come up in either of the following circumstances:

- When the neighbor address (numbered or unnumbered) has a /32 mask.
- When the unnumbered local address for ANCP dynamic logical interfaces is configured to use a preferred source address.

ANCP neighbors have one of the following administrative states, which simply represent the configuration of the neighbor:

- **enabled**—The neighbor is configured in the CLI.
- **disabled**—The neighbor is not configured, meaning either that it has never been configured or that the configuration has been deleted.

ANCP neighbors in the enabled state have one of the following operational states, which represent the state of adjacency negotiations:

- **Configured**—The neighbor has been configured, but has never established an adjacency.
- **Establishing**—Adjacency negotiations are in progress.
- **Established**—Adjacency negotiations have succeeded and an ANCP session has been established.
- **Not Established**—The neighbor has lost a previously established adjacency, but is ready to begin negotiations.

You can also configure parameters for a specific neighbor that override global or default configurations by including any of the following statements at the `[edit protocols ancp neighbor ip-address]` hierarchy level:

- **adjacency-timer**—Adjust the interval between adjacency messages exchanged with this neighbor.
- **ietf-mode**—Prevent the ANCP agent from operating in a backward-compatible mode for this neighbor; for neighbors that use the current IETF implementation of ANCP.
- **maximum-discovery-table-entries**—Specify how many discovery table entries are accepted from this neighbor. Include this statement at the `[edit protocols ancp]` hierarchy level to set the number of entries globally for all neighbors.
- **pre-ietf-mode**—Enable the ANCP agent to operate in a backward-compatible mode for this neighbor; for neighbors that use the original IETF implementation of ANCP (GSMPv2) rather than the current

implementation. Include this statement at the [edit protocols ancp] hierarchy level to operate in backward-compatible mode globally for all neighbors.

RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks, defines ANCP Version 1. ANCP was originally implemented based on General Switch Management Protocol (GSMP) version 3, sub-version 1. However, the Internet community has made so many extensions and modifications to GSMPv3 in the course of developing ANCP that ANCP is no longer interoperable with GSMPv3. Consequently, ANCP neighbors must be able to dynamically detect the version that each peer supports. A joint registry codifies the GSMP and ANCP version numbers.

When an ANCP neighbor opens adjacency negotiations, it indicates the highest version of ANCP that it supports, either 0x31 for GSMPv3 or 0x32 for ANCP Version 1. (Version 1 may also be called Version 50, referring to the decimal conversion from the hexadecimal value.) If the receiving neighbor supports that version of ANCP, it returns that value when it responds to the sending neighbors. If it does not support that version, the receiving neighbor simply drops the message.

The ANCP agent stores information about active ANCP subscribers in the Junos shared database, including DSL attributes for the access lines. This storage is persistent and is removed from the database only when you delete the interface or interface set for the access line or issue one of the following commands:

- *clear ancp neighbor*
- *clear ancp subscriber*

The persistence of the storage enables PPPoE and DHCP IP demux subscribers to be properly managed by RADIUS for authentication and accounting, with their DSL attributes, even when the ANCP connection has been temporarily terminated.

Partitions

ANCP supports the division of an access node into logical partitions. Each partition creates an adjacency with a router; each partition on an access node can form adjacencies with different routers. Partition negotiation takes place during ANCP adjacency negotiation. ANCP messages carry the following fields relating to the partition negotiation:

- The partition type (PType) field indicates whether the access node is partitioned and how the partition identifier is negotiated. The field has one of the following values negotiated during the formation of the adjacency:
 - 0—The access node is not partitioned or does not support partitions.
 - 1—The number of partitions is fixed and the router requests the access node to use the identifier it places in the partition identifier field.
 - 2—The number of partitions is fixed and the access node has assigned the partition identifier.

- The partition ID field that indicates one of the following scenarios for ANCP agent support of the neighbor:
 - Zero partition ID—The ANCP agent supports each neighbor on an IP address over a single TCP session with a partition ID of zero. This is the default support case. This value is required when the partition type is zero.
 - Single nonzero partition ID—The ANCP agent supports each neighbor on an IP address over a single TCP session with a nonzero partition ID. This case requires partition ID learning to be enabled with the `gsmpp-syn-wait` statement at the `[edit protocols ancp]` hierarchy level.
- The partition flag (PFlag) field indicates the type of partition request being made. A value of one specifies a new adjacency.

The following partitioning schemes are supported

- Each partition has an independent ANCP session and channel to an adjacent router. All partitions have a fixed partition ID of zero.
- Each partition has an independent ANCP session and channel to an adjacent router. Each partition has a dedicated, nonzero partition ID.

Adjacency Update Messages

After an adjacency has been established, the ANCP agent uses adjacency update messages to inform routers that control the same partition about each other. Once more than one router has established an adjacency to a given partition, the ANCP agent sends an adjacency update message to each of these routers to report how many established adjacencies the partition currently supports. When an adjacency is lost, an update message is sent to the remaining routers to report the change in status. You can use the `show ancp neighbor detail` command to display the number of adjacencies currently established on a particular partition.

Generic Response Messages and Result Codes

ANCP neighbors and the router can reply to messages either with a specific response message or a generic response message. A generic response message is typically sent when no information needs to be sent to the peer other than a success or failure result. If the response is about a failure, then a result code is included that specifies the kind of failure; a limited amount of diagnostic data can also be included. A generic response message can also be sent independently of a request if the adjacency is being shut down because of the failure. In this case, the sender of the message zeros out the Transaction ID field in the message header and the Message Type field in the Status-Info TLV.

[Table 64 on page 1630](#) describes the result codes that can be included in a generic response message.

Table 64: ANCP Failure Result Codes

Code Value	Description	Detected By
0x02	Although the request message is properly formed, it is invalid because it violates the protocol, either because of timing issues such as a race condition or the direction in which the message was transmitted.	ANCP agent
0x06	One or more of the specified ports is down because of a state mismatch between the router and an ANCP control application.	Control applications (none yet available)
0x13	ANCP is out of resources. This result code is sent only by the access node; the problem is probably not related to the access lines, but can be related to a specific request.	ANCP protocol layer or control applications (none yet available)
0x51	The type of request message is not implemented because of a mismatch in protocol versions or capability state between the peers, or possibly because the message type is optional for an ANCP capability.	ANCP agent
0x53	The message is malformed either because it was corrupted in transit or an implementation error occurred at one end of the connection.	ANCP agent
0x54	One or more mandatory TLVs is missing from the request.	ANCP agent

Table 64: ANCP Failure Result Codes *(Continued)*

Code Value	Description	Detected By
0x55	The contents of one or more TLVs in the request are invalid because they do not match the TLV specification.	ANCP agent
0x500	One or more of the ports specified in a request does not exist, possibly because of a configuration mismatch between the access node and the router or AAA.	Control applications (none yet available)



NOTE: Although Junos OS supports both sending and receiving generic response messages, currently the ANCP agent only receives these messages. When one of these messages is received, the router generates a system log, increments the generic message counters, and increments the result code counters. When the ANCP agent receives an incorrect or unexpected generic response message from an ANCP neighbor, it immediately drops the packet, generates a system log notice message, and takes no further action.

Generic response messages usually include the Status-Info TLV, which includes supplemental information about a warning or error condition. The Status-Info TLV is required when the result code indicates any of the following: a port is down or does not exist, a mandatory TLV is missing, or a TLV is invalid. The Status-Info TLV can also be included in other ANCP message types.

Benefits of Access Node Control Protocol

- Simplify the configuration and maintenance of access lines between access nodes and subscribers.
- Perform CoS-related adjustments on upstream and downstream data rate attributes to both accurately provide services and control congestion in the network.
- Provide access network information, such as DSL attributes to backend applications such as operations support systems (OSS) for service management.
- Store DSL attributes in the session database for use during RADIUS authentication and accounting of PPPoE sessions.

SEE ALSO

| *Agent Circuit Identifier-Based Dynamic VLANs Overview*

ANCP Operations in Different Network Configurations**IN THIS SECTION**

- [1:1 and N:1 Traffic Shaping Models | 1633](#)
- [Business Services Traffic Shaping Model | 1635](#)
- [ANCP Network Using N:1 and 1:1 Configuration Models without Interface Sets | 1636](#)
- [Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets | 1637](#)
- [ANCP Network Using N:1 Configuration Model with Interface Sets | 1638](#)
- [Sequence of ANCP Events: Static VLAN Interfaces over Ethernet with Interface Sets | 1640](#)
- [ANCP Network Using 1:1 Configuration Model with Interface Sets | 1640](#)
- [Sequence of ANCP Events: Static VLAN Demux Interfaces over Aggregated Ethernet with Interface Sets | 1642](#)

This topic describes different types of supported network configurations and the sequence of events for ANCP operations in representative sample network topologies.

You can configure the ANCP agent for any of the following interface types:

- Static VLAN interfaces, except those created by Extensible Subscriber Services Manager (ESSM)
- Static VLAN demux interfaces, except those created by ESSM
- Static interface sets, including those created by ESSM
- Dynamic interface sets
- Dynamic VLAN-tagged interface sets
- Dynamic agent circuit identifier (ACI) interface sets, also known as ACI sets or ACI VLANs
- Dynamic PPPoE and DHCP IP demux subscriber interfaces
- Dynamic VLAN demux interfaces with Ethernet-VPLS encapsulation

Subscriber sessions are dynamically created as needed for each of the devices in a household. Each household can include multiple CPE devices that access the Internet. In all cases, each household is

identified by a unique ACI that is assigned by the access node. Additional identifiers are used in some configurations.

1:1 and N:1 Traffic Shaping Models

The 1:1 and N:1 traffic shaping models determine how VLANs are correlated with households. These models are also referred to as access models or configuration models. A network can include one or both of the models:

- **1:1 model**—A household has only one PPPoE or DHCP IP demux subscriber session. One or more such households can exist on a single VLAN or VLAN demux interface. In the case of a single household, either the subscriber interface or its underlying VLAN or VLAN demux interface can represent the household. In the case of multiple households, the corresponding subscriber interfaces represent the households. In either case, the interface representing a household must be mapped to the ACI for its access line.

[Table 65 on page 1633](#) describes the types of interfaces supported for the ANCP 1:1 access model when interface sets are not involved, and whether the PPPoE or DHCP IP demux discovery packets must include the ACI for the subscriber access lines.

Table 65: ACI Mapping by Interface Type for the ANCP 1:1 Model

Interface Type	Description	Presence of ACI in Discovery Packets
Dynamic PPPoE or DHCP IP demux interface	When ACI is present in discovery packets, the ANCP agent maps the ACI to the subscriber interface. The name of the interface is automatically generated and nondeterministic.	Required.
Static VLAN or VLAN demux interface	The name of the interface is statically configured. The ANCP agent configuration must include the access-identifier statement to statically map the ACI to the interface.	Not present.

- **N:1 model**—A household can have more than one PPPoE or DHCP IP demux subscriber session. The household can have more than one VLAN or VLAN demux interface. In either case, all the interfaces must be grouped into an interface set. The interface set in turn must be mapped to the ACI for the household's access line.

An interface set groups the dynamic PPPoE or DHCP IP demux sessions for a household. The subscribers are placed into interface sets by one several methods. [Table 66 on page 1634](#) describes the types of interface sets supported in the ANCP N:1 access model, how they are created, and how the ACI is mapped to the interface set.

Table 66: ACI Mapping by Interface Set Type for the ANCP N:1 Access Model

Type of Interface Set	Description	Interface Type	Presence of ACI in Discovery Packets
ACI-based VLAN interface sets	<p>When the router receives a DHCP or PPPoE discovery packet that includes an ACI embedded within the DSL Forum vendor-specific tag, it dynamically creates the VLAN and the interface set. The router generates a nondeterministic name for the interface set, such as aci-1003-ge-1/0/0.1073741832.</p> <p>The ANCP agent automatically maps the ACI from the discovery packet to the dynamically created interface set.</p> <p>All DHCP IP demux or PPPoE sessions that have the same ACI are mapped to the same interface set.</p>	Dynamic VLAN and VLAN demux interfaces.	Required.
Dynamic interface sets	<p>A dynamic profile dynamically creates the interface set and places interfaces in the set. The profile can either have the name of the interface set explicitly configured or a variable that represents the interface set name. If a variable is used, then the interface set name is provided by RADIUS when it returns an Access-Accept message for the subscriber.</p> <p>The ANCP agent configuration must include the access-identifier statement to statically map the ACI to the interface set.</p> <p>All DHCP IP demux and PPPoE sessions are mapped to an interface set according to the rules of the dynamic profile.</p>	DHCP IP demux subscriber interfaces, PPPoE subscriber interfaces, or VLAN interfaces.	Irrelevant.
Static interface sets	<p>The interface set and set name are statically configured and include multiple static interfaces.</p> <p>The ANCP agent configuration must include the access-identifier statement to statically map the ACI to the interface set.</p>	Static VLAN and VLAN demux interfaces.	Irrelevant.

Table 66: ACI Mapping by Interface Set Type for the ANCP N:1 Access Model (Continued)

Type of Interface Set	Description	Interface Type	Presence of ACI in Discovery Packets
VLAN-tagged interface sets	<p>When the router receives a DHCP or PPPoE discovery packet that includes a VLAN ID, it dynamically creates the VLAN and the interface set. The interface set is given a deterministic name consisting of the physical interface name and the VLAN tags, for example, ge-1/0/0-101.</p> <p>The ANCP agent configuration must include the access-identifier statement to statically map the ACI to the interface set.</p> <p>All DHCP IP demux or PPPoE sessions that have the same VLAN ID tag are mapped to the same interface set.</p>	Dynamic VLAN and VLAN demux interfaces.	Irrelevant.

CoS traffic shaping is based on the subscriber downstream traffic rate that the ANCP agent receives from the access node and then passes to CoS. CoS can shape subscriber traffic at the level of the household or the session:

- Household shaping—Only aggregate traffic to the household is shaped. Household shaping results from applying a CoS traffic-control profile to the static VLAN or VLAN demux interface or to the interface set.
- Session shaping—The traffic rate to individual devices in the household is shaped. Session shaping results from specifying a CoS traffic-control profile in the dynamic PPPoE profile that creates the subscriber session. Depending on the network configuration, session shaping may employ shared priority queues to shape all sessions identically or individual priority queues to shape the sessions separately.

Business Services Traffic Shaping Model

In addition to the N:1 and 1:1 traffic shaping models, the ANCP agent also supports a business services traffic shaping model. In this model, the Extensible Subscriber Services Manager (ESSM) classifies a PPPoE session as either residential household or business subscriber. Classification occurs during RADIUS authentication and authorization. The ANCP agent applies CoS traffic shaping differently depending on the classification.

Before RADIUS authentication and authorization, the PPPoE session represents a residential household in the ANCP 1:1 model. The ANCP agent dynamically maps the household's access line to the

corresponding subscriber interface and applies CoS traffic shaping to that interface. The household line's ACI is present in the PPPoE discovery packet.

When ESSMD subsequently classifies a PPPoE session as a business subscriber session during RADIUS authentication and authorization, it creates and groups multiple management and data plane static VLAN interfaces into a static interface set. then it statically maps the access line for the PPPoE session to this interface set according to the CLI configuration. The ANCP agent removes CoS traffic shaping from the subscriber interface and applies it to the static interface set. Removing the CoS traffic shaping means that the CoS application applies the next rate in its default or configured adjustment control profile to the interface or interface set. The new business subscriber interface set cannot contain a mix of static and dynamic interfaces. That prohibition is not limited to dynamic VLANs and the PPPoE session that triggered the creation of the interface set.



NOTE: An exception to the ANCP agent's general support for CoS traffic shaping and RADIUS authentication and accounting on static VLAN and VLAN Demux interfaces is that it does not support these interfaces if they are created by ESSM. These interfaces are different from the ESSM-created interface sets, which are supported by the ANCP agent.


From the perspective of the ANCP agent, the business services model effectively overrides a dynamically derived access-line-to-interface mapping with a statically configured access-line-to-interface-set mapping. This action triggers the ANCP agent to reapply CoS traffic shaping accordingly.

The business services model is typically used in a Layer 2 wholesale network. For detailed information, see *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*.

ANCP Network Using N:1 and 1:1 Configuration Models without Interface Sets

In this sample topology, two households are configured for one underlying static VLAN or VLAN demux interface (N:1; dual-tagged VLAN) and a single household is configured for another underlying interface (1:1; single-tagged VLAN) ([Figure 47 on page 1637](#)). In addition to the unique ACI assigned by the access node, each household is further identified by the VLAN, which is mapped to the identifier in the ANCP agent configuration. CoS traffic shaping for sessions can employ only shared priority queues to shape all sessions identically; individual priority queues to shape the sessions separately are not supported.

Figure 47: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model)



Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets

The following sequence of events is for the topology in [Figure 47 on page 1637](#) with static VLAN interfaces over Ethernet without interface sets.

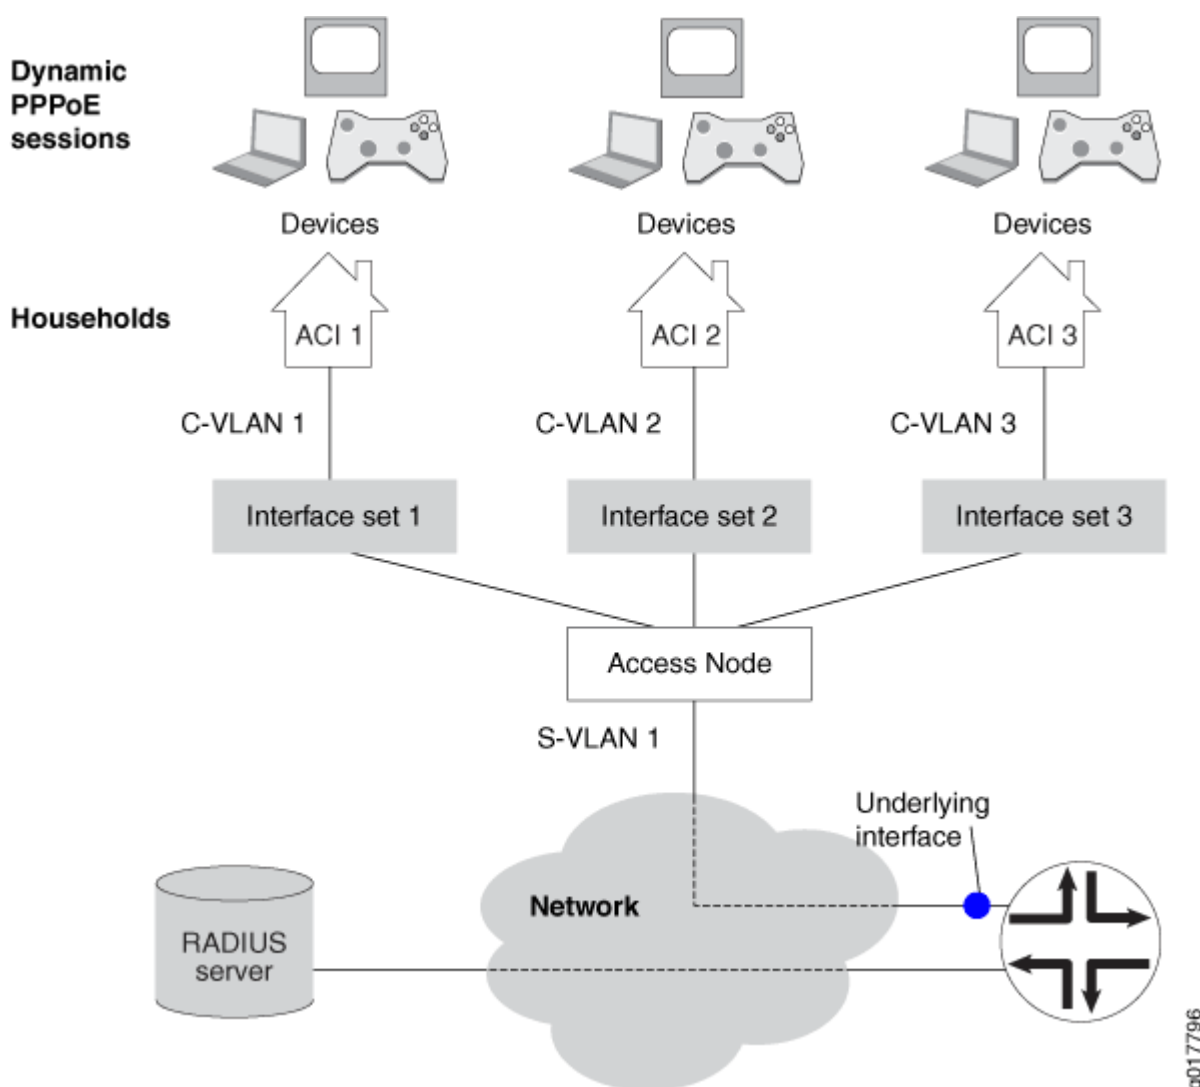
1. A network device in the household initiates PPPoE discovery.
2. PPPoE creates a dynamic PPPoE session on the underlying static VLAN or VLAN demux interface and applies the advisory options configured on the VLAN to the session.
3. The access node independently provides the ANCP agent with the ANCP DSL attributes for an access line identified by an ACI.
4. The ANCP agent sends CoS the adjusted downstream data rate for the static VLAN or demux VLAN mapped to the ACI. The ANCP agent stores all DSL attributes, including the adjusted upstream data rate, in the router's shared database.
5. AAA correlates the dynamic PPPoE session with the access line by matching the underlying interface of the session to the static VLAN or VLAN demux interface associated with the ACI in the ANCP agent configuration.

6. AAA retrieves the ANCP DSL attributes for the access line from the router's shared database and maps them to the Juniper Networks DSL VSAs in the RADIUS Access-Request and Accounting-Request messages. If the DSL attributes are unavailable, the session's advisory upstream and downstream data rates are mapped to the Upstream-Calculated-Qos-Rate VSA (26-142) and Downstream-Calculated-Qos-Rate (26-141) VSAs, respectively. These VSAs are then included in the RADIUS messages.

ANCP Network Using N:1 Configuration Model with Interface Sets

In this topology, multiple households are configured for each underlying static VLAN or VLAN demux interface ([Figure 48 on page 1639](#)). The VLANs are dual-tagged. Each household includes several CPE devices that access the Internet. In addition to the unique ACI assigned by the access node, the household is further identified by the interface set. The interface set groups the dynamic PPPoE sessions for the individual subscriber devices. It is either explicitly configured in the dynamic PPPoE profile or specified in the RADIUS Access-Accept message during PPPoE session authentication. Session shaping can employ shared priority queues to shape all sessions identically or individual queues to shape the sessions separately.

Figure 48: Sample ANCP Topology with Interface Sets (N:1 Model)



In this N:1 model with interface sets, the access node must add the DSL Forum VSA to the PPPoE PADI and PADR discovery packets that it passes to the router during the establishment of dynamic PPPoE sessions. The VSA includes the ACI for the household. This inclusion enables AAA to correlate the PPPoE sessions with their respective subscriber access lines and DSL attributes during RADIUS authentication and accounting. If the ACI is not present, AAA cannot make the correlation and subsequently reports only the advisory upstream and downstream data rates to RADIUS Authentication and Accounting.

When the dynamic PPPoE profile is configured with the `$junos-interface-set-name` predefined variable, the configuration of the access node, router, and RADIUS server must be synchronized with regard to the ACI and interface set:

- The RADIUS Access-Accept message must contain the Juniper Networks Qos-Interface-Set-Name VSA (26-130).

- The CoS Layer 2 configuration must explicitly identify the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).
- The ANCP agent configuration must map an ACI to the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).

Sequence of ANCP Events: Static VLAN Interfaces over Ethernet with Interface Sets

The following sequence of events is for the topology in [Figure 48 on page 1639](#) with static VLAN interfaces over Ethernet with interface sets.


1. A network device in the household initiates PPPoE discovery.
2. The access node adds the DSL Forum VSA tag with the ACI for the household to the PPPoE PADI and PADR discovery packets. (The identifier is known to PPPoE as the agent circuit identifier.)
3. PPPoE creates a dynamic PPPoE session with the provided ACI on the underlying static VLAN and applies the advisory options configured on the VLAN to the session.
4. The access node independently provides the ANCP agent with the ANCP DSL attributes for an access line identified by an ACI.
5. The ANCP agent provides CoS with the adjusted downstream data rate for the interface set mapped to the ACI. The ANCP agent stores all ANCP DSL attributes, including the adjusted upstream and downstream data rates, in the router's shared database.
6. AAA correlates the dynamic PPPoE session with the access line by matching the session identifier received in the DSL Forum VSA to the ACI configured for the interface set in the ANCP agent configuration.
7. AAA retrieves the ANCP DSL attributes for the access line from the router's shared database and maps them to the Juniper Networks DSL VSAs in the RADIUS Access-Request and Accounting-Request messages. If the DSL attributes are unavailable, the session's advisory upstream and downstream data rates are mapped to the Upstream-Calculated-Qos-Rate VSA (26-142) and Downstream-Calculated-Qos-Rate (26-141) VSAs, respectively. These VSAs are then included in the RADIUS messages.
8. When authentication is completed, the dynamic PPPoE session is placed into the interface set configured in the dynamic PPPoE profile. The profile specifies a named interface set or the `$junos-interface-set-name` predefined variable, which indicates that the interface set is named in the RADIUS Access-Accept message.

ANCP Network Using 1:1 Configuration Model with Interface Sets

In this topology, a single household is configured for each underlying static VLAN or VLAN demux interface ([Figure 49 on page 1641](#)). The VLANs are dual-tagged. Each household includes several CPE

devices that access the Internet. In addition to the unique ACI assigned by the access node, the household is further identified by the interface set. The interface set is either explicitly configured in the dynamic PPPoE profile or specified in the RADIUS Access-Accept message during PPPoE session authentication. Session shaping can employ shared priority queues to shape all sessions identically or individual queues to shape the sessions separately.

Figure 49: Sample ANCP Topology with Interface Sets (1:1 Model)



In this 1:1 model with interface sets, the ANCP agent configuration must map the underlying interface for the PPPoE sessions in an interface set to both the ACI and the interface set. This configuration enables AAA to correlate the PPPoE sessions with their respective subscriber access lines and DSL attributes during RADIUS authentication and accounting.

When the dynamic PPPoE profile is configured with the `$junos-interface-set-name` predefined variable, the configuration of the access node, router, and RADIUS server must be synchronized with regard to the ACI and interface set:

- The RADIUS Access-Accept message must contain the Juniper Networks Qos-Interface-Set-Name VSA (26-130).

- The CoS Layer 2 configuration must explicitly identify the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).
- The ANCP agent configuration must map an ACI to the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).

Sequence of ANCP Events: Static VLAN Demux Interfaces over Aggregated Ethernet with Interface Sets

The following sequence of events is for the topology in [Figure 49 on page 1641](#) with static VLAN demux interfaces over aggregated Ethernet with interface sets.

1. A network device in the household initiates PPPoE discovery.
2. PPPoE creates a dynamic PPPoE session with the provided ACI on the underlying static VLAN demux interface and applies the advisory options configured on the VLAN to the session.
3. The access node independently provides the ANCP agent with the ANCP DSL attributes for an access line identified by an ACI.
4. The ANCP agent provides CoS with the adjusted downstream data rate for the interface set mapped to the ACI. The ANCP agent stores all ANCP DSL attributes, including the adjusted upstream and downstream data rates, in the router's shared database.
5. AAA correlates the dynamic PPPoE session with the access line by matching the underlying interface of the session to the underlying interface configured for the interface set in the ANCP agent configuration.
6. AAA retrieves the ANCP DSL attributes for the access line from the router's shared database and maps them to the Juniper Networks DSL VSAs in the RADIUS Access-Request and Accounting-Request messages. If the DSL attributes are unavailable, the session's advisory upstream and downstream data rates are mapped to the Upstream-Calculated-Qos-Rate VSA (26-142) and Downstream-Calculated-Qos-Rate (26-141) VSAs, respectively. These VSAs are then included in the RADIUS messages.
7. When authentication is completed, the dynamic PPPoE session is placed into the interface set configured in the dynamic PPPoE profile. The profile specifies a named interface set or the `$junos-interface-set-name` predefined variable, which indicates that the interface set is named in the RADIUS Access-Accept message.

Configuring the ANCP Agent

You can configure the ANCP agent to enable a service-oriented Layer 3 edge device to discover information about the topology of a connected access network. The ANCP agent can also provide details about subscriber traffic and enable the adjustment of QoS traffic shaping for subscribers.

To configure the ANCP agent:

1. Specify each ANCP neighboring access node to be monitored and optionally configure neighbor parameters.
See ["Configuring ANCP Neighbors" on page 1644.](#)
2. Specify the subscribers reached by a VLAN or a set of VLANs through a particular access node.
See ["Associating an Access Node with Subscribers for ANCP Agent Operations" on page 1645.](#)
3. (Optional) Configure the adjacency timer.
See ["Specifying the Interval Between ANCP Adjacency Messages" on page 1646.](#)
4. (Optional) Specify the maximum number of discovery table entries that are accepted.
See ["Specifying the Maximum Number of Discovery Table Entries" on page 1647](#)
5. (Optional) Configure the ANCP agent to work with an early IETF draft.
See ["Configuring the ANCP Agent for Backward Compatibility" on page 1647.](#)
6. (Optional) Configure the graceful restart timer.
See ["Specifying How Long Processes Wait for the ANCP Agent Restart to Complete" on page 1648.](#)
7. (Optional) Configure the ANCP agent to learn partition IDs from neighbors.
See ["Configuring the ANCP Agent to Learn ANCP Partition IDs" on page 1649.](#)
8. (Optional) Configure an adjustment factor per DSL line type for the downstream and upstream data rates that the ANCP agent reports to AAA.
See ["Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates" on page 1695.](#)
9. (Optional) Configure an adjustment factor per PON line type for the downstream and upstream data rates that the ANCP agent reports to AAA.
See ["Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates" on page 1697.](#)
10. (Optional) Configure the ANCP agent to report unadjusted downstream traffic rates to CoS.
See ["Configuring the ANCP Agent to Report Traffic Rates to CoS" on page 1688.](#)
11. (Optional) Specify a recommended shaping rate to be applied by RADIUS to downstream or upstream traffic per ANCP interface.
See ["Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces" on page 1693.](#)
12. (Optional) Configure AAA to Include or Exclude Juniper Networks access line VSAs in RADIUS authentication and accounting messages.
See ["Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages" on page 1713.](#)
13. (Optional) Configure AAA to send an immediate interim accounting update to the RADIUS server when AAA receives a rate change notification from the ANCP agent on the router.
See ["Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications" on page 1714.](#)

14. (Optional) Configure the ANCP agent to associate a neighbor with an access-facing physical interface for the creation of autosensed dynamic VLANs on the interface.

See *Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs*.

15. (Optional) Configure the ANCP agent to dampen the effect of short-term adjacency losses for all neighbors.

See *Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses*.

16. (Optional) Configure the ANCP agent to dynamically generate interface set names for business subscribers.

See *How to Configure the Automatic Creation of Business Subscriber Interface Sets*.

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

See ["Tracing ANCP Events for Troubleshooting" on page 1722](#).

Configuring ANCP Neighbors

You must configure each neighboring access node that you want the ANCP agent to monitor and potentially shape traffic for. Some neighbor settings override globally configured values.

To configure an ANCP neighbor:

1. Specify the IP address of the neighbor.

```
[edit protocols ancp]
user@host# set neighbor 203.0.113.234
```

2. (Optional) Configure the neighbor to operate in a backward-compatible mode when it does not support the current IETF standard and the backward-compatible mode is not configured globally.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set pre-ietf-mode
```

3. (Optional) Override the globally configured backward-compatible mode when the neighbor supports the current IETF standard.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set ietf-mode
```

4. (Optional) Configure the interval in seconds between ANCP adjacency messages exchanged with this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set adjacency-timer 20
```

5. (Optional) Specify the maximum number of discovery table entries that are accepted from this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set maximum-discovery-table-entries 10000
```

6. (Optional) Enable out-of-band ANCP triggering of autosensed, dynamic VLANs on the physical interface.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set auto-configure-trigger interface ge-1/0/0
```

7. (Optional) Configure how long the ANCP agent maintains a Layer 2 wholesale session when an adjacency loss occurs.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set adjacency-loss-hold-time 10
```

Associating an Access Node with Subscribers for ANCP Agent Operations

The ANCP agent on the router uses the access loop circuit identifier (ACI) to distinguish individual ANCP subscribers. Because the agent uses the ACI to associate (map) each subscriber to an interface or interface set, each ACI must be unique across all ANCP neighbors connected to the router.



BEST PRACTICE: We recommend that the ACIs be unique across your ANCP network.

Access lines can be statically or dynamically mapped to interfaces or interface set. When the subscriber's DHCP or PPPoE discovery packets contain the ACI, then the agent can dynamically map it to the interface or interface set. Otherwise, the ACI must be statically configured. A static configuration overrides dynamic mapping of ACIs—and therefore subscribers—to interfaces or sets.

You can use the `access-identifier` statement only for interface and interface set types that have configured or deterministic names: static VLAN interfaces, static VLAN demux interfaces, static interface sets, dynamic interface sets, and dynamic VLAN-tagged interface sets.

The `access-identifier` statement is required for mapping an access line to static interface sets, dynamic interface sets, and dynamic VLAN-tagged interface sets. This is true regardless of the presence of an ACI in the PPPoE or DHCP IP demux subscriber's discovery packet, because the use of the ACI is irrelevant to the creation of these types of interface sets.

You cannot use the `access-identifier` statement for the following interface and interface set types, because they have nondeterministic, automatically generated names: dynamic VLAN demux interfaces, dynamic ACI interface sets (ACI VLANs), and dynamic PPPoE and DHCP IP demux subscriber interfaces.

To associate an ACI with a set of VLAN interfaces for subscribers:

- Specify the name of the interface set and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set interface-set vlan5 access-identifier "dslam port 2/3"
```

To associate an ACI with a single VLAN:

- Specify the logical interface and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set ge-1/0/4.12 access-identifier "dslam port-2-10"
```

To associate an ACI with a static VLAN demux interface:

- Specify the logical interface and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set demux0.100 access-identifier aci_100_1_0
```

Specifying the Interval Between ANCP Adjacency Messages

When the ANCP agent and a neighbor negotiate to establish an adjacency, each proposes a value for the interval between the adjacency messages that they exchange after it is established. The larger of the values proposed by the agent and the neighbor is selected for the interval between subsequent adjacency messages exchanged by the agent and the neighbor. You can specify the interval value that the ANCP agent proposes for either all neighbors or a specific neighbor.

To configure the proposed interval between ANCP adjacency messages for all neighbors:

- Specify the time in seconds.

```
[edit protocols ancp]
user@host# set adjacency-timer 20
```

To configure the proposed interval between ANCP adjacency messages for a specific neighbor:

- Specify the time in seconds.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set adjacency-timer 20
```

Specifying the Maximum Number of Discovery Table Entries

You can specify the maximum number of discovery table entries accepted from all neighbors or from a particular neighbor.

To configure the maximum number of entries for all neighbors:

- Specify the number of entries.

```
[edit protocols ancp]
user@host# set maximum-discovery-table-entries 5000
```

To configure the maximum number of entries for a specific neighbor:

- Specify the number of entries.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set maximum-discovery-table-entries 5000
```

Configuring the ANCP Agent for Backward Compatibility

You can configure the ANCP agent to operate in a mode compatible with the protocol as it was initially proposed to operate. This backward-compatible or pre-IETF mode is compatible with Internet draft draft-wadhwa-gsmp-l2control-configuration-00.txt, *GSMP extensions for layer2 control (L2C)*. Setting this backward-compatible mode enables interoperability with devices that are not compatible with the later ANCP Internet drafts or RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*.

When this mode is configured globally for all neighbors, you can override it for a particular neighbor that supports the IETF draft or standard.

To configure the ANCP agent to operate in a backward-compatible mode for all neighbors:

- Specify the pre-IETF mode.

```
[edit protocols ancp]
user@host# set pre-ietf-mode
```

To configure the ANCP agent to operate in a backward-compatible mode for a specific neighbor:

- Specify the pre-IETF mode.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set pre-ietf-mode
```

- To override the globally configured backward-compatible mode for a specific neighbor:

Specify the IETF mode.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set ietf-mode
```

Specifying How Long Processes Wait for the ANCP Agent Restart to Complete

You can specify how long other processes wait for the ANCP agent to restart. The ANCP agent sends a keepalive message to CoS at intervals equal to one-third the value of the maximum helper restart time. For example, when you configure the maximum restart time to 120 seconds, the ANCP agent sends a keepalive message every 40 seconds.

If CoS does not receive a keepalive message within the maximum helper restart time, it considers the ANCP agent to be down and immediately reverts any traffic shaping updates that were implemented as a result of ANCP agent monitoring to the configured values. Consequently, traffic to the subscribers is not effectively shaped, potentially resulting in traffic drops in the DSLAMs. The configured values are maintained until the ANCP agent comes back up and sends fresh traffic shaping updates to CoS.

To configure how long other processes wait for the ANCP agent to restart:

- Specify the time in seconds.

```
[edit protocols ancp]
user@host# set maximum-helper-restart-time 150
```

Configuring the ANCP Agent to Learn ANCP Partition IDs

By default, the ANCP agent expects ANCP partition IDs to be zero, meaning that the access node is not divided into logical partitions that can each form adjacencies with routers. You can configure the ANCP agent to support nonzero partition IDs. When you do so, the agent waits a configurable period to receive a SYN message from a neighbor during adjacency initiation. When the agent receives such a message, it uses the partition information contained in the Partition ID, PType, and PFlag fields to generate in turn a SYN message that it sends to the neighbor to continue adjacency negotiation.

To configure the ANCP agent to learn partition ID information from neighbors:

1. Enable partition ID learning.

```
[edit protocols ancp]
user@host# set gsmp-syn-wait
```

2. (Optional) Specify the maximum time the ANCP agent waits to receive a SYN message from a neighbor during the formation of an adjacency.

```
[edit protocols ancp]
user@host# set gsmp-syn-timeout seconds
```

For example, to enable partition ID learning and force the ANCP agent to wait 45 seconds for a SYN message:

```
[edit protocols ancp]
user@host# set gsmp-syn-wait
user@host# set gsmp-syn-timeout 45
```


Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet

IN THIS SECTION

- [Requirements | 1650](#)
- [Overview | 1651](#)
- [Configuration | 1658](#)
- [Verification | 1676](#)

This example describes how to configure an ANCP network topology that manages subscriber access for several households by grouping individual devices into interface sets, providing access and services through one dedicated C-VLAN per household, and shaping traffic on a per-household basis. In this N:1 configuration, dual-tagged VLANs are configured over a single, underlying, static VLAN demux interfaces over aggregated Ethernet.

Requirements

This example uses the following hardware and software components:

- MX Series 5G Universal Routing Platform with only MPCs installed for VLAN demux support
- RADIUS server
- DSLAM access node

Before you begin configuring the example, be sure you have:

- Thoroughly read and understood the following topics:
 - ["ANCP and the ANCP Agent Overview" on page 1622](#)
 - ["ANCP Operations in Different Network Configurations" on page 1632](#)
- Configured your access node.
- Configured your RADIUS server.

Overview

IN THIS SECTION

- [Topology | 1652](#)

ANCP provides a means to configure, maintain, and monitor local access lines between access nodes (DSLAMs) and subscribers. Associated CoS configurations shape the downstream subscriber traffic. ANCP can enable more accurate traffic shaping by adjusting net data rates to discount the packet overhead of the access lines and then providing these adjusted rates to CoS.

The network topology in this example includes a dual-tagged (C-VLAN/S-VLAN) VLAN configuration over a static VLAN demux interface that is in turn configured over aggregated Ethernet for redundancy. This topology is an N:1 configuration model because—although each C-VLAN corresponds to one subscriber household—all the C-VLANs are configured over the same underlying VLAN demux interface. Multiple end-user devices in each household—or rather the dynamic PPPoE sessions established by each device—are grouped by household into interface sets. The grouping is accomplished by a separate dynamic profile configured for each C-VLAN. The ANCP agent configuration maps the ACI for the household's access line to an interface set. CoS applies a traffic-control profile to each interface set to shape the subscriber-directed traffic on a per-household basis. The CoS shaping rate is dynamically updated based upon the DSL attributes provided by the access node for each household's access line.

[Figure 50 on page 1652](#) shows S-VLAN 103, configured on demux0, servicing the access node. C-VLANs 1, 2, and 3 each service a single household (subscriber). The respective households are identified by unique ACIs. The dynamic PPPoE sessions for devices in each household are grouped for monitoring and traffic shaping into interface sets 10301, 10302, and 10303.

Topology

Figure 50: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet

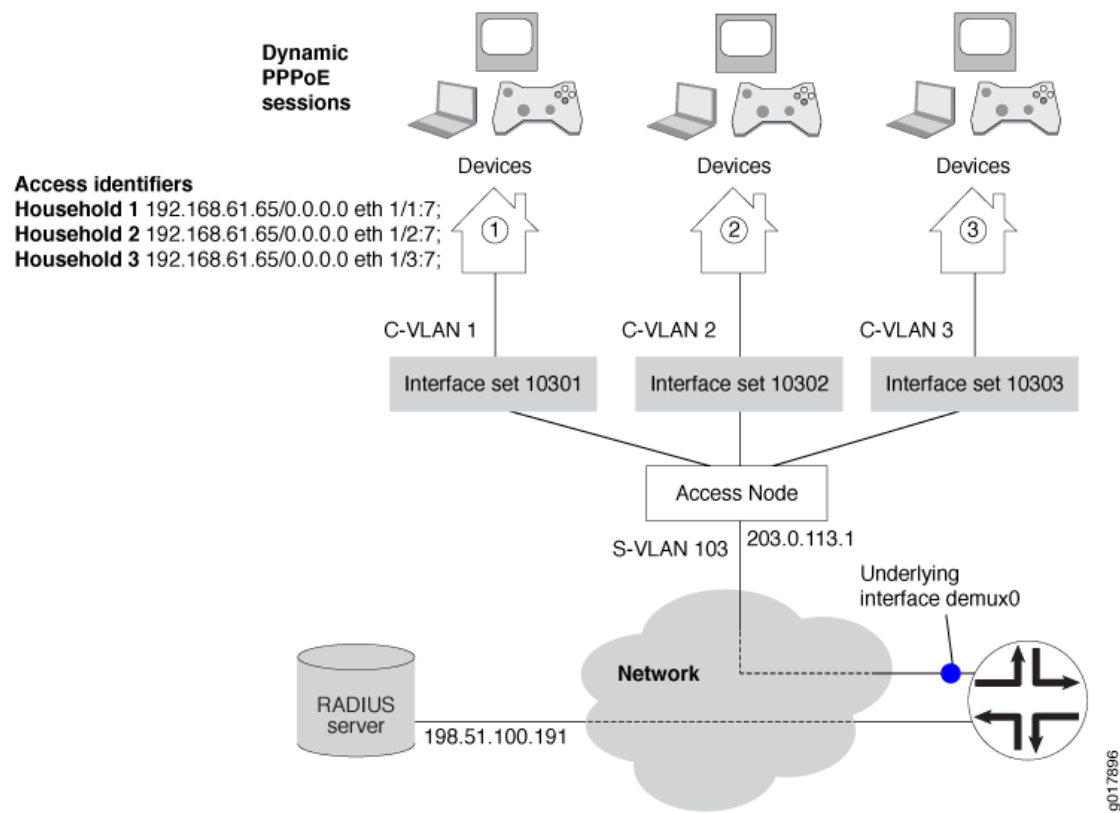


Table 67 on page 1653 describes the configuration components used in this example.

Table 67: Configuration Components used in ANCP N:1 Topology Example with Interface Sets

Configuration Component or Property	Component Name or Setting	Description
Dynamic profiles	ancp-10301	<p>Each profile defines the dynamic PPPoE session created when any of the devices for a particular subscriber household accesses the network.</p> <p>Each profile specifies the following:</p> <ul style="list-style-type: none"> • A set of interfaces in which the sessions are created. • Dynamic instantiation of both the logical interfaces for the sessions and the underlying PPPoE logical interfaces on which the subscribers log in. • CHAP and PAP authentication for the sessions. • The interval between successive PPP keepalive messages. • The loopback address for the dynamic PPPoE logical interfaces.
	ancp-10302	
	ancp-10303	
Predefined variables	\$junos-interface-unit	Instantiates the logical interface for each PPPoE session.
	\$junos-underlying-interface	Instantiates the logical underlying PPP interface on which each dynamic PPPoE logical interface is created when a subscriber logs in.

Table 67: Configuration Components used in ANCP N:1 Topology Example with Interface Sets
(Continued)

Configuration Component or Property	Component Name or Setting	Description
Interfaces	ae0	<p>Aggregated Ethernet interface that is the underlying interface for the VLAN demux interfaces.</p> <p>The interface includes the following configuration:</p> <ul style="list-style-type: none"> • CoS hierarchical scheduling. • Stacked VLAN tagging for all logical interfaces on top of ae0. • Link protection.
	demux0	VLAN demux interface that runs over the underlying aggregated Ethernet interface.

Table 67: Configuration Components used in ANCP N:1 Topology Example with Interface Sets
(Continued)

Configuration Component or Property	Component Name or Setting	Description
	demux0.10301 demux0.10302 demux0.10303	<p>VLAN demux logical interfaces that correspond to the C-VLANs for individual subscriber households.</p> <p>Each logical interface includes the following configuration:</p> <ul style="list-style-type: none"> • Inner (C-VLAN) and outer VLAN (S-VLAN) tags. • The underlying physical interface, ae0. • The dynamic profile that creates PPPoE sessions on the C-VLAN. • Downstream and upstream advisory traffic rates. • Proxy ARP and protection against duplicate sessions on the interface.
	ge-1/0/1	Primary member link in the aggregated Ethernet bundle.
	ge-1/0/2	Backup member link in the aggregated Ethernet bundle.
	lo0.0	Loopback interface for use in the access network. The loopback interface is automatically used for unnumbered interfaces.

Table 67: Configuration Components used in ANCP N:1 Topology Example with Interface Sets
(Continued)

Configuration Component or Property	Component Name or Setting	Description
	pp0	PPP interface on which the PPPoE subscriber logical interfaces are created.
Interface sets	10301 10302 10303	Set of interfaces in which the sessions for the devices in a particular household are created. Each interface set is specified in a dynamic profile for that household. ANCP associates each interface set with an ACI and a VLAN demux logical interface (C-VLAN). CoS applies a traffic-control profile to each interface set.
Advisory traffic rates	downstream-rate	Recommended rate for downstream traffic in the absence of traffic rate information from the access node.
	upstream-rate	Recommended rate for upstream traffic in the absence of traffic rate information from the access node.
Traffic-control profile	tcp1	CoS profile that shapes the downstream subscriber traffic rate; in this example, shaping is adjusted for ATM packet overhead. The profile is applied to the interface sets.

Table 67: Configuration Components used in ANCP N:1 Topology Example with Interface Sets
(Continued)

Configuration Component or Property	Component Name or Setting	Description
IP addresses	203.0.113.1	Address of the ANCP access node that monitors the subscriber households.
	127.0.50.1/28	Address of the loopback interface, lo0.
	198.51.100.191	Address of the RADIUS accounting server and authentication server.
Access circuit loop identifiers	192.168.61.65/0.0.0.0 eth 1/1:7; 192.168.61.65/0.0.0.0 eth 1/2:7; 192.168.61.65/0.0.0.0 eth 1/3:7;	Identifier for the local access circuit from the access node to the subscriber household. It identifies the household. ANCP associates each identifier with an interface set.

The ANCP agent configuration includes the following elements:

- The IP address for the access node (DSLAM) is specified as 203.0.113.1. The interval between ANCP adjacency messages sent between neighbors is set to 5 seconds.
- The ANCP agent is enabled to report adjusted data rates to CoS to improve the accuracy of downstream traffic shaping. The ANCP agent adjusts the net data rates for ADSL lines by ninety percent and for ADSL2 lines by ninety-five percent.
- Each interface set is associated with both the ACI unique to the subscriber household and the relevant underlying VLAN demux interface.

The RADIUS configuration on the router includes the following elements:

- The IP address (198.51.100.191) for the authentication and accounting server, as well as the secret password for accessing the server.
- The subscriber access profile, radius-profile, specifies that RADIUS is used for authentication.

- Juniper Networks DSL VSAs are included in RADIUS request messages, but the DSL Forum VSA attributes are excluded from RADIUS messages
- Accounting sessions are configured to be recognized in decimal format.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1658](#)
- [Configuring the Dynamic PPPoE Profiles | 1661](#)
- [Configuring the Static VLAN Demux Interface over Aggregated Ethernet | 1664](#)
- [Configuring Class of Service | 1670](#)
- [Configuring ANCP | 1672](#)
- [Configuring RADIUS Authentication and Accounting | 1674](#)

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 ANCP network with static N:1 demux VLANs to the subscriber households, perform these tasks:

CLI Quick Configuration

To quickly configure the ANCP network described in this example, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
# Dynamic Profiles
edit dynamic-profiles ancp-10301
set interfaces interface-set 10301 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
set ppp-options chap
set ppp-options pap
set pppoe-options underlying-interface "$junos-underlying-interface"
set keepalives interval 30
set family inet unnumbered-address lo0.0
top
edit dynamic-profiles ancp-10302
```

```

set interfaces interface-set 10302 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
set ppp-options chap
set ppp-options pap
set pppoe-options underlying-interface "$junos-underlying-interface"
set keepalives interval 30
set family inet unnumbered-address lo0.0
top
edit dynamic-profiles ancp-10303
set interfaces interface-set 10303 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
set ppp-options chap
set ppp-options pap
set pppoe-options underlying-interface "$junos-underlying-interface"
set keepalives interval 30
set family inet unnumbered-address lo0.0
top
#
# Aggregated Ethernet Interfaces and VLAN Demux Interfaces
set interfaces ge-1/0/1 hierarchical-scheduler
set interfaces ge-1/0/1 gigether-options 802.3ad ae0
set interfaces ge-1/0/1 gigether-options 802.3ad primary
set interfaces ge-1/0/2 hierarchical-scheduler
set interfaces ge-1/0/2 gigether-options 802.3ad ae0
set interfaces ge-1/0/2 gigether-options 802.3ad backup
set interfaces ae0 hierarchical-scheduler
set interfaces ae0 stacked-vlan-tagging
set interfaces ae0 aggregated-ether-options link-protection
set interfaces demux0 unit 10301 proxy-arp
set interfaces demux0 unit 10301 vlan-tags outer 103
set interfaces demux0 unit 10301 vlan-tags inner 1
set interfaces demux0 unit 10301 demux-options underlying-interface ae0
set interfaces demux0 unit 10301 family pppoe duplicate-protection
set interfaces demux0 unit 10301 family pppoe dynamic-profile ancp-10301
set interfaces demux0 unit 10301 advisory-options downstream-rate 16m
set interfaces demux0 unit 10301 advisory-options upstream-rate 1m
set interfaces demux0 unit 10302 proxy-arp
set interfaces demux0 unit 10302 vlan-tags outer 103
set interfaces demux0 unit 10302 vlan-tags inner 2
set interfaces demux0 unit 10302 demux-options underlying-interface ae0
set interfaces demux0 unit 10302 family pppoe duplicate-protection
set interfaces demux0 unit 10302 family pppoe dynamic-profile ancp-10302
set interfaces demux0 unit 10302 advisory-options downstream-rate 16m

```

```

set interfaces demux0 unit 10302 advisory-options upstream-rate 1m
set interfaces demux0 unit 10303 proxy-arp
set interfaces demux0 unit 10303 vlan-tags outer 103
set interfaces demux0 unit 10303 vlan-tags inner 3
set interfaces demux0 unit 10303 demux-options underlying-interface ae0
set interfaces demux0 unit 10303 family pppoe duplicate-protection
set interfaces demux0 unit 10303 family pppoe dynamic-profile ancp-10303
set interfaces demux0 unit 10303 advisory-options downstream-rate 16m
set interfaces demux0 unit 10303 advisory-options upstream-rate 1m
set interfaces lo0 unit 0 family inet address 127.0.50.1/28
top
#
# Class of Service
edit class-of-service
set traffic-control-profiles tcp1 shaping-rate 16m
set traffic-control-profiles tcp1 overhead-accounting cell-mode
set interfaces interface-set 10301 output-traffic-control-profile tcp1
set interfaces interface-set 10302 output-traffic-control-profile tcp1
set interfaces interface-set 10303 output-traffic-control-profile tcp1
top
#
# ANCP
edit protocols ancp
set traceoptions file ancpd
set traceoptions file size 512m
set traceoptions flag config
set traceoptions flag cos
set qos-adjust
set adjacency-timer 5
set maximum-helper-restart-time 90
set qos-adjust-adsl 90
set qos-adjust-adsl2 95
set interfaces interface-set 10301 access-identifier "192.168.61.65/0.0.0.0 eth 1/1:7;"
set interfaces interface-set 10302 access-identifier "192.168.61.65/0.0.0.0 eth 1/2:7;"
set interfaces interface-set 10303 access-identifier "192.168.61.65/0.0.0.0 eth 1/3:7;"
set interfaces interface-set 10301 underlying-interface demux0.10301
set interfaces interface-set 10302 underlying-interface demux0.10302
set interfaces interface-set 10303 underlying-interface demux0.10303
set neighbor 203.0.113.1
top
#
# RADIUS
edit access

```

```

set radius-server 198.51.100.191 secret "$ABC123$ABC123$ABC123"
edit access profile radius-profile
set authentication-order radius
set radius authentication-server 198.51.100.191
set radius accounting-server 198.51.100.191
set radius options accounting-session-id-format decimal
set radius options juniper-dsl-attributes
set radius attributes exclude dsl-forum-attributes access-request
set radius attributes exclude dsl-forum-attributes accounting-start
set radius attributes exclude dsl-forum-attributes accounting-stop
top

```

Configuring the Dynamic PPPoE Profiles

Step-by-Step Procedure

In this procedure, you configure a dynamic profile for each C-VLAN: ancp-10301, ancp-10302, and ancp-1033.

1. Configure the interface set that the PPPoE sessions on this C-VLAN are placed in.

```

[edit dynamic-profiles ancp-10301]
user@host1# edit interfaces interface-set 10301

```

2. Configure the logical interfaces to be dynamically instantiated for the interface set.

```

[edit dynamic-profiles ancp-10301 interfaces interface-set 10301]
user@host1# set interface pp0 unit "$junos-interface-unit"

```

3. Configure CHAP and PAP authentication as properties of the dynamic PPPoE logical interfaces.

```

[edit dynamic-profiles ancp-10301 interfaces pp0 unit "$junos-interface-unit"]
user@host1# set ppp-options chap
user@host1# set ppp-options pap

```

4. Configure the logical underlying interface on which the router creates the dynamic PPPoE logical interface; this is the interface on which the subscriber logs in.

```
[edit dynamic-profiles ancp-10301 interfaces pp0 unit "$junos-interface-unit"]
user@host1# set pppoe-options underlying-interface "$junos-underlying-interface"
```

5. Specify the interval between successive keepalive requests.

```
[edit dynamic-profiles ancp-10301 interfaces pp0 unit "$junos-interface-unit"]
user@host1# set keepalives interval 30
```

6. Configure the IPv4 protocol family and that the local (unnumbered) address can be derived from the loopback address for the dynamic PPPoE logical interfaces.

```
[edit dynamic-profiles ancp-10301 interfaces pp0 unit "$junos-interface-unit"]
user@host1# set family inet unnumbered-address lo0.0
```

7. Repeat Steps 1 through 6 for the second dynamic profile, ancp-10302, and the third dynamic profile, ancp-10303.

Results

From configuration mode, confirm the dynamic profile configuration by entering the `show dynamic-profiles` command.

```
[edit]
user@host# show dynamic-profiles
ancp-10301 {
  interfaces {
    interface-set 10301 {
      interface pp0 {
        unit "$junos-interface-unit";
      }
    }
  }
  pp0 {
    unit "$junos-interface-unit" {
      ppp-options {
        chap;
        pap;
      }
    }
  }
}
```



```

pp0 {
  unit "$junos-interface-unit" {
    ppp-options {
      chap;
      pap;
    }
    pppoe-options {
      underlying-interface "$junos-underlying-interface";
    }
    keepalives interval 30;
    family inet {
      unnumbered-address lo0.0;
    }
  }
}
}
}
}

```

When you are done configuring the device, enter `commit` from configuration mode.

Configuring the Static VLAN Demux Interface over Aggregated Ethernet

Step-by-Step Procedure

1. Enable hierarchical scheduling on this interface.

```

[edit interfaces ge-1/0/1]
user@host1# set hierarchical-scheduler

```

2. Specify this interface as the primary member of the aggregated Ethernet bundle.

```

[edit interfaces ge-1/0/1]
user@host1# set together-options 802.3ad ae0 primary

```

3. Enable hierarchical scheduling on a second interface.

```

[edit interfaces ge-1/0/2]
user@host1# set hierarchical-scheduler

```

4. Specify this interface as the backup member of the aggregated Ethernet bundle.

```
[edit interfaces ge-1/0/2]
user@host1# set gigether-options 802.3ad ae0 backup
```

5. Enable hierarchical scheduling on the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set hierarchical-scheduler
```

6. Enable stacked VLAN tagging for all logical interfaces on the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set stacked-vlan-tagging
```

7. Enable link protection as a property of the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set aggregated-ether-options link-protection
```

8. Configure VLAN demux interface demux0.10301.

- a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set proxy-arp
```

- b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10301]
user@host1# set vlan tags outer 103 inner 1
```


- c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set demux-options underlying-interface ae0
```

- d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set family pppoe duplicate-protection
```

- e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set family pppoe dynamic-profile ancp-10301
```

- f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10301]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

9. Configure VLAN demux interface demux0.10302.

- a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set proxy-arp
```

- b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10302]
user@host1# set vlan tags outer 103 inner 2
```

- c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set demux-options underlying-interface ae0
```

- d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set family pppoe duplicate-protection
```

- e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set family pppoe dynamic-profile ancp-10302
```

- f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10302]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

10. Configure VLAN demux interface demux0.10303.

- a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set proxy-arp
```

- b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10303]
user@host1# set vlan tags outer 103 inner 3
```

- c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set demux-options underlying-interface ae0
```

- d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set family pppoe duplicate-protection
```

- e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set family pppoe dynamic-profile ancp-10303
```

- f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10303]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

11. Configure the IPv4 protocol family and the address of the loopback interface.

```
[edit interfaces lo0]
user@host1# set unit 0 family inet address 127.0.50.1/28
```

Results

From configuration mode, confirm the static VLAN demux configuration by entering the `show interfaces` command.

```
[edit]
user@host# show interfaces
ge-1/0/1 {
  hierarchical-scheduler;
  together-options {
```

```

        802.3ad {
            ae0;
            primary;
        }
    }
}
ge-1/0/2 {
    hierarchical-scheduler;
    gigether-options {
        802.3ad {
            ae0;
            backup;
        }
    }
}
ae0 {
    hierarchical-scheduler;
    stacked-vlan-tagging;
    aggregated-ether-options {
        link-protection;
    }
}
demux0 {
    unit 10301 {
        proxy-arp;
        vlan-tags outer 103 inner 1;
        demux-options {
            underlying-interface ae0;
        }
        family pppoe {
            duplicate-protection;
            dynamic-profile ancp-10301;
        }
        advisory-options {
            downstream-rate 16m;
            upstream-rate 1m;
        }
    }
    unit 10302 {
        proxy-arp;
        vlan-tags outer 103 inner 2;
        demux-options {
            underlying-interface ae0;
        }
    }
}

```

```

    }
    family pppoe {
        duplicate-protection;
        dynamic-profile ancp-10302;
    }
    advisory-options {
        downstream-rate 16m;
        upstream-rate 1m;
    }
}
unit 10303 {
    proxy-arp;
    vlan-tags outer 103 inner 3;
    demux-options {
        underlying-interface ae0;
    }
    family pppoe {
        duplicate-protection;
        dynamic-profile ancp-10303;
    }
    advisory-options {
        downstream-rate 16m;
        upstream-rate 1m;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.50.1/28
        }
    }
}
}

```

When you are done configuring the device, enter `commit` from configuration mode.

Configuring Class of Service

Step-by-Step Procedure

1. Configure the traffic-control profile with the shaping rate and specify the overhead accounting mode to account for ATM cell encapsulation.

```
[edit class-of-service]
user@host1# set traffic-control-profiles tcp1 shaping-rate 16m
user@host1# set traffic-control-profiles tcp1 overhead-accounting cell-mode
```

2. Apply the traffic-control profile to the interface sets.

```
[edit class-of-service]
user@host1# set interfaces interface-set 10301 output-traffic-control-profile tcp1
user@host1# set interfaces interface-set 10302 output-traffic-control-profile tcp1
user@host1# set interfaces interface-set 10303 output-traffic-control-profile tcp1
```

Results

From configuration mode, confirm the class of service configuration by entering the `show class-of-service` command.

```
[edit]
user@host# show class-of-service
traffic-control-profiles {
    tcp1 {
        shaping-rate 16m;
        overhead-accounting cell-mode;
    }
}
interfaces {
    interface-set 10301 {
        output-traffic-control-profile tcp1;
    }
    interface-set 10302 {
        output-traffic-control-profile tcp1;
    }
    interface-set 10303 {
        output-traffic-control-profile tcp1;
    }
}
```

When you are done configuring the device, enter `commit` from configuration mode.

Configuring ANCP

Step-by-Step Procedure

1. Configure the access node address.

```
[edit protocols ancp]
user@host1# set neighbor 203.0.113.1
```

2. Configure the ANCP agent to report adjusted downstream traffic rates to CoS.

```
[edit protocols ancp]
user@host1# set qos-adjust
```

3. Specify an overhead adjustment of the traffic on ADSL and ADSL2 lines to 90 percent and 95 percent, respectively, of the net data rate.

```
[edit protocols ancp]
user@host1# set qos-adjust-ads1 90
user@host1# set qos-adjust-ads12 95
```

4. Specify an interval of 5 seconds between adjacency messages sent to all ANCP neighbors.

```
[edit protocols ancp]
user@host1# set adjacency-timer 5
```

5. Associate the ACI with the interface sets for each C-VLAN.

```
[edit protocols ancp]
user@host1# set interfaces interface-set 10301 access-identifier "192.168.61.65/0.0.0.0 eth
1/1:7;"
user@host1# set interfaces interface-set 10302 access-identifier "192.168.61.65/0.0.0.0 eth
1/2:7;"
```

```
user@host1# set interfaces interface-set 10303 access-identifier "192.168.61.65/0.0.0.0 eth
1/3:7;"
```

6. Specify the underlying interface for the interface sets.

```
[edit protocols ancp]
user@host1# set interfaces interface-set 10301 underlying-interface demux0.10301
user@host1# set interfaces interface-set 10302 underlying-interface demux0.10302
user@host1# set interfaces interface-set 10303 underlying-interface demux0.10303
```

7. Configure the size of the ANCP trace log files.

```
[edit protocols ancp traceoptions]
user@host1# set file ancpd size 512m
```

8. Configure flags for tracing ANCP configuration and CoS operations.

```
[edit protocols ancp traceoptions]
user@host1# set flag config
user@host1# set flag cos
```

Results

From configuration mode, confirm the ANCP agent configuration by entering the `show ancp` command.

```
[edit]
user@host# show ancp
traceoptions {
  file ancpd size 512m;
  flag config;
  flag cos;
}
qos-adjust;
adjacency-timer 5;
qos-adjust-adsl 90;
qos-adjust-adsl2 95;
interfaces {
  interface-set {
```



```

    10301 {
        access-identifier "192.168.61.65/0.0.0.0 eth 1/1:7;";
        underlying-interface demux0.10301;
    }
    10302 {
        access-identifier "192.168.61.65/0.0.0.0 eth 1/2:7;";
        underlying-interface demux0.10302;
    }
    10303 {
        access-identifier "192.168.61.65/0.0.0.0 eth 1/3:7;";
        underlying-interface demux0.10303;
    }
}
neighbor 203.0.113.1;

```

When you are done configuring the device, enter `commit` from configuration mode.

Configuring RADIUS Authentication and Accounting

Step-by-Step Procedure

1. Configure the password for the RADIUS server.

```

[edit access]
user@host1# set radius-server 198.51.100.191 secret "$ABC123$ABC123$ABC123"

```

2. Specify that RADIUS is used to authenticate subscribers.

```

[edit access]
user@host1# set profile radius-profile authentication-order radius

```

3. Configure the RADIUS authentication and accounting server.

```

[edit access]
user@host1# set profile radius-profile radius authentication-server 198.51.100.191
user@host1# set profile radius-profile radius accounting-server 198.51.100.191

```

4. Configure options for the RADIUS server: The format used to identify the accounting session and that Juniper Networks DSL VSAs are added to RADIUS request messages.

```
[edit access]
user@host1# set profile radius-profile radius options accounting-session-id-format decimal
user@host1# set profile radius-profile radius options juniper-dsl-attributes
```

5. Exclude DSL Forum VSA attributes from being included in RADIUS messages.

```
[edit access]
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes access-
request
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes
accounting-start
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes
accounting-stop
```

Results

From configuration mode, confirm the RADIUS configuration by entering the `show access` command.

```
[edit]
user@host# show access
radius-server {
    198.51.100.191 secret "$ABC123$ABC123$ABC123"; ## SECRET-DATA
}
profile radius-profile {
    radius {
        authentication-server 198.51.100.191;
        accounting-server 198.51.100.191;
        options {
            accounting-session-id-format decimal;
            juniper-dsl-attributes;
        }
        attributes {
            exclude {
                dsl-forum-attributes [ access-request accounting-start accounting-stop ];
            }
        }
    }
}
```

```
}
}
```

When you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Aggregated Ethernet Interface Configuration | 1676](#)
- [Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set | 1677](#)
- [Verifying the demux0 Interface Configuration | 1678](#)
- [Verifying the pp0 Interface Configuration | 1679](#)
- [Verifying the ANCP Agent Configuration | 1680](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying the Aggregated Ethernet Interface Configuration

Purpose

Verify that the interface values match your configuration, the link is up, and traffic is flowing.

Action

From operational mode, enter the `show interfaces redundancy` command.

```
user@host> show interfaces redundancy
```

Interface	State	Last change	Primary	Secondary	Current status
ae0	On primary		ge-1/0/1	ge-1/0/2	both up

From operational mode, enter the `show interfaces ae0` command.

```
user@host> show interfaces ae0
```

Physical interface: ae0, Enabled, Physical link is Up

```

Interface index: 128, SNMP ifIndex: 606
Link-level type: Ethernet, MTU: 1522, Speed: 1Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
Last flapped   : 2012-03-11 13:24:18 PST (2d 03:34 ago)
Input rate      : 1984 bps (2 pps)
Output rate     : 0 bps (0 pps)

```

```

Logical interface ae0.32767 (Index 69) (SNMP ifIndex 709)
Flags: SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :          371259          2    46036116    1984
  Output:              0           0         0         0
Protocol multiservice, MTU: Unlimited
Flags: Is-Primary

```

Meaning

The `show interfaces redundancy` output shows the redundant link configuration and that both link interfaces are up. The `show interfaces ae0` output shows that the aggregated Ethernet interface is up and that traffic is being received on the logical interface.

Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set

Purpose

Verify that the traffic scheduling and shaping parameters are configured and applied properly.

Action

```
user@host> show class-of-service
```

Verifying the demux0 Interface Configuration

Purpose

Verify that the VLAN demux interface displays the configured PPPoE family attributes and the member links in the aggregated Ethernet bundle.

Action

From operational mode, enter the `show interfaces demux0` command for each VLAN.

```
user@host> show interfaces demux0.10301
Logical interface demux0.10301 (Index 76) (SNMP ifIndex 61160)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
  Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 199)
  Link:
    ge-1/0/1
    ge-1/0/2
  Input packets : 2
  Output packets: 18575
  Protocol pppoe
    Dynamic Profile: ancp-10301,
    Service Name Table: None,
    Max Sessions: 16000, Duplicate Protection: On,
    AC Name: pppoe-server-1
```

Alternatively, you can enter `show pppoe underlying-interfaces detail` to display the state and PPPoE family configuration for all configured underlying interfaces.

Meaning

The output shows the name of the underlying interface, the member links of the aggregated bundle, and the PPPoE family configuration. The output shows packet counts when traffic is present on the logical interface.

Verifying the pp0 Interface Configuration

Purpose

Verify that the interface values match your configuration.

Action

From operational mode, enter the `show interfaces pp0` command.

```
user@host> show interfaces pp0.100
Logical interface pp0.100 (Index 71) (SNMP ifIndex 710)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
PPPoE:
  State: SessionUp, Session ID: 1,
  Session AC name: pppoe-server-1, Remote MAC address: 00:00:5E:00:53:34,
  Underlying interface: demux0.10301 (Index 70)
Link:
  ge-5/0/3.32767
  ge-5/1/2.32767
Input packets : 18572
Output packets: 18572
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 18566 (00:00:02 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Success
  Protocol inet, MTU: 1500
  Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Is-Primary
    Local: 203.0.113.45
```

Meaning

This output shows information about the PPPoE logical interface created on the underlying VLAN demux interface. The output includes the PPPoE family and aggregated Ethernet redundant link information, and shows input and output traffic for the PPPoE interface.

Verifying the ANCP Agent Configuration

Purpose

Verify that the ANCP values match your configuration and that traffic is flowing.

Action

From operational mode, enter the `show ancp subscriber` command.

```
user@host> show ancp subscriber detail
```

Interface	State	Last change	Primary	Secondary	Current status
ae0	On primary		ge-1/0/1	ge-1/0/2	both up

From operational mode, enter the `show ancp cos` command.

```
user@host> show ancp cos
```

```

Qos Adjust Flag:      TRUE
Keepalive Timer:      30 secs
Cos State:             WRITE_READY
Connect Time:         Mon Mar 19 15:03:01 2012
Session Time:         Mon Mar 19 15:03:13 2012
Routing Instance Time: Mon Mar 19 15:03:14 2012
Keepalive Time:       Not Set
Rate Update Time:     Mon Mar 19 15:03:15 2012

```

Type	Name	Index	Pending Update	Last Update
iflset	10301	1	None	64 Kbps
iflset	10302	2	None	64 Kbps
iflset	10303	71	None	64 Kbps

Meaning

The `show ancp subscriber` output shows subscriber line information such as state and the various traffic rates collected by the ANCP agent—displayed for each subscriber as identified by the ACI. The `show ancp cos` output shows that the ANCP agent is configured to send adjusted rate data to CoS, that keepalives

are configured for a 30-second interval, and that the interface sets 10301, 10302, and 10303 are configured and their traffic rates are updating

SEE ALSO

Dynamic Profiles Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Subscriber Interfaces and Demultiplexing Overview

[ANCP Agent Interactions with AAA | 1701](#)

[ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | 1703](#)

[Configuring the ANCP Agent | 1642](#)

[AAA Service Framework Overview | 2](#)

RELATED DOCUMENTATION

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

ANCP Agent Traffic Shaping and CoS

IN THIS SECTION

- [Traffic Rate Reporting and Adjustment by the ANCP Agent | 1682](#)
- [Preservation of CoS Shaping Across ANCP Agent Restarts | 1687](#)
- [Configuring the ANCP Agent to Report Traffic Rates to CoS | 1688](#)
- [Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | 1693](#)
- [Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1695](#)
- [Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | 1697](#)
- [Verifying and Monitoring CoS for ANCP Subscribers | 1699](#)

Traffic Rate Reporting and Adjustment by the ANCP Agent

IN THIS SECTION

- [Overview | 1682](#)
- [Traffic Rate Adjustment | 1684](#)
- [Recommended Traffic Shaping Rates | 1686](#)
- [ANCP Agent Keepalives for CoS | 1687](#)

The ANCP agent monitors the subscriber access lines and reports to AAA and CoS information about the lines that it receives from the access node. Starting in Junos OS Release 17.4R1, the ANCP agent can use access line information that it receives in the PPPoE intermediate agent (PPPoE-IA) tags. This information is carried in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x91) in PADI or PADR packets. In earlier releases, the ANCP agent can use only access line information that it receives in ANCP messages. The access line information for both carriers is logically sourced from the same data on the access node; it represents a current, accurate snapshot of the values at the moment that the subscriber connection is initiated.

It is theoretically possible for ANCP and PPPoE subscribers to specify different data rates in the Vendor-Specific-Tags TLV when the connection is first established. This is an unlikely occurrence, but when the dynamic profile is configured to accept these values, the most recently received value takes precedence. The rates announced on the PPPoE connection are expected to be used only when ANCP is either not used or does not include rates. However, network dynamics make it impossible to guarantee the source from which the information arrives first. If the values conflict, a subsequent Port Up message from the access node forces the resolution to the ANCP values.

Overview

The ANCP agent reports two kinds of data rates:

- The *net data rate* is the portion of the total data rate that can be used to transmit user information. The net data rate is also called the *unadjusted* traffic rate.
- Because each DSL line type has a certain technology overhead, the actual rate for user data is less than the net data rate. The *adjusted* or *calculated* rate is the net data rate reduced by the amount of technology overhead incurred by each DSL line type. The result is a closer approximation of the actual rate of subscriber data traffic.

The ANCP agent reports traffic rates differently to AAA and CoS:

- The agent always reports unadjusted rates for both upstream and downstream traffic to AAA in response to a AAA request. When configured, the agent adjusts the traffic rates and reports the adjusted values in addition to the unadjusted rate.
- The ANCP agent reports traffic rates to CoS only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level. The agent reports only downstream traffic rates to CoS in support of CoS traffic shaping. It never reports upstream traffic rates to CoS because CoS does not shape upstream traffic. The agent also reports to CoS the overhead mode and bytes for the access line; CoS can use this information when it subsequently shapes the traffic.

When you remove a shaping rate configuration that the ANCP agent previously applied, the traffic shaping rate reverts to the CoS session shaping as determined by the CoS traffic-control profiles specified in the dynamic profile. However, if the ANCP agent remains running but loses a connection to a particular neighbor whose subscriber traffic has been adjusted as a result of ANCP agent action, the adjusted rate remains in effect. The rate currently in effect changes only when the ANCP agent restores the connection and sends fresh updates to CoS, or when you remove the `qos-adjust` statement.

Because CoS can perform traffic shaping only when a traffic-control profile has been applied to the interface or interface set, you might expect the ANCP agent to always influence traffic shaping when the ANCP subscriber interface or interface set has a traffic-control profile. This behavior does not always occur.

Consider a configuration where a subscriber logical interface is a member of an ACI-based VLAN (interface set); all members share the same ACI. The dynamic profile that instantiates the subscriber interface applies a traffic-control profile to the interface. The profile that instantiates the VLAN applies an interface-shared filter instead of a traffic-control profile.

The following sequence of events takes place when the subscriber logs in.

1. The first packet creates the auto-sensed, underlying VLAN.
2. The second packet creates the ACI-based subscriber VLAN
3. The third packet creates the subscriber logical interface.

Because the VLAN comes up first, the ANCP agent attaches to the VLAN and not to the interface. Consequently, the agent reports to CoS the downstream data rate only for the VLAN, not for the logical interface. CoS has no information to adjust the shaping rate for the interface, so it shapes traffic for the interface only according to the interface's traffic-control profile.

Although the agent does report the downstream rate for the VLAN, CoS cannot use that information to shape the VLAN traffic, because the VLAN does not have a traffic-control profile. Consequently, the VLAN rate does not affect the logical interface's rate even though the logical interface is a member of that interface set.

Traffic Rate Adjustment

When a DSLAM determines the data rate on the subscriber local loop, it ignores the additional headers on the DSL line that are associated with the overhead of the access mode (ATM or Ethernet) and the technology of the DSL line type. When the ANCP agent subsequently reports a net data rate, by default it includes this overhead, reporting a slightly higher value than the actual subscriber data rate seen by the DSLAM.

You can configure the ANCP agent to additionally report an adjusted rate to account for the traffic overhead. The ANCP agent dynamically adjusts the net data rate by applying a fixed percentage value to the net data rate received from the DSLAM. The percentage adjustment factor applies globally for all subscribers of the particular DSL line type as follows:

- The agent can adjust the rates it reports to AAA for all DSL types.
- The agent can adjust the rates it reports to CoS for only frame-mode DSL types (SDSL, VDSL, VDSL2, and OTHER). It cannot adjust the rates reported to CoS for cell-mode DSL types (ADSL, ADSL2, and ADSL2+).

You can also configure the ANCP agent to adjust the number of overhead bytes that it reports to CoS per cell or frame. The agent can add or subtract the specified value from the actual number of overhead bytes for all DSL types. The agent does not report the number of overhead bytes (adjusted or unadjusted) to AAA.

[Table 68 on page 1684](#) summarizes how adjusted rates and overheads are reported.

Table 68: Traffic Adjustment Reporting by Access Line Type

DSL Access Line Type	Upstream and Downstream Adjusted Rate Reported to AAA	Downstream-Only Adjusted Rate Reported to CoS	Adjusted Overhead Byte Count Reported to CoS
ADSL	✓	–	✓
ADSL2	✓	–	✓
ADSL2+	✓	–	✓
OTHER	✓	✓	✓
SDSL	✓	✓	✓

Table 68: Traffic Adjustment Reporting by Access Line Type (Continued)

DSL Access Line Type	Upstream and Downstream Adjusted Rate Reported to AAA	Downstream-Only Adjusted Rate Reported to CoS	Adjusted Overhead Byte Count Reported to CoS
VDSL	✓	✓	✓
VDSL2	✓	✓	✓

The ANCP agent reports traffic rates to CoS only when you have included the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level. By default, these are unadjusted rates. CoS attempts to avoid traffic drops in the access node by itself adjusting the traffic shaping rate that it applies to downstream traffic for a particular VLAN or set of VLANs. The discrepancy between the actual user data rate and the agent-reported net data rate reduces the accuracy of CoS traffic shaping. You increase the accuracy of CoS traffic shaping by configuring the ANCP agent to report adjusted rate and byte values to CoS.

If you are running Junos OS Release 17.3 or earlier, use the CLI configuration statements in [Table 69 on page 1685](#) to make traffic adjustments. The CoS statements are at the `[edit protocols ancp qos-adjust]` hierarchy level. The AAA statements are at the `[edit protocols ancp]` hierarchy level.

Table 69: Statements for Adjusting Traffic Rate and Overhead by Access Line Type Through Junos OS Release 17.3

Access Line Type	Adjust Net Downstream and Upstream Rates for AAA	Adjust Net Downstream Rates for CoS	Adjust Overhead Bytes for CoS
ADSL	<code>qos-adjust-adsl</code>	–	<code>adsl-bytes</code>
ADSL2	<code>qos-adjust-adsl2</code>	–	<code>adsl2-bytes</code>
ADSL2+	<code>qos-adjust-adsl2-plus</code>	–	<code>adsl2-plus-bytes</code>
OTHER	<code>qos-adjust-other</code>	<code>other-overhead-adjust</code>	<code>other-bytes</code>
SDSL	<code>qos-adjust-sdsl</code>	<code>sdsl-overhead-adjust</code>	<code>sdsl-bytes</code>
VDSL	<code>qos-adjust-vdsl</code>	<code>vdsl-overhead-adjust</code>	<code>vdsl-bytes</code>

Table 69: Statements for Adjusting Traffic Rate and Overhead by Access Line Type Through Junos OS Release 17.3 (Continued)

Access Line Type	Adjust Net Downstream and Upstream Rates for AAA	Adjust Net Downstream Rates for CoS	Adjust Overhead Bytes for CoS
VDSL2	qos-adjust-vdsl2	vdsl2-overhead-adjust	vdsl2-bytes

If you are running Junos OS Release 17.4R1 or later, use the `access-line` configuration statement options in [Table 70 on page 1686](#) to make traffic adjustments for CoS and AAA options. The `access-line` statement is at the `[edit system]` hierarchy level.

Table 70: Statements for Adjusting Traffic Rate and Overhead by Access Line Type Starting in Junos OS Release 17.4R1

Access Line Type	Adjust Net Downstream and Upstream Rates for AAA	Adjust Net Downstream Rates for CoS	Adjust Overhead Bytes for CoS
ADSL	adsl-total-adjust	–	adsl-overhead-bytes
ADSL2	adsl2-total-adjust	–	adsl2-overhead-bytes
ADSL2+	adsl2-plus-total-adjust	–	adsl2-plus-overhead-bytes
OTHER	other-total-adjust	other-overhead-adjust	other-overhead-bytes
SDSL	sdsl-total-adjust	sdsl-overhead-adjust	sdsl-overhead-bytes
VDSL	vdsl-total-adjust	vdsl-overhead-adjust	vdsl-overhead-bytes
VDSL2	vdsl-total-adjust	vdsl2-overhead-adjust	vdsl2-overhead-bytes

Recommended Traffic Shaping Rates

To handle a situation where the router does not receive information from the access node about the downstream and upstream calculated traffic rates for an interface, you can specify recommended *advisory* values for shaping the traffic sent to the interface so that it matches the subscriber local loop speed.

The transmit speed is the recommended traffic value in bits per second used for downstream traffic for an ANCP interface, and is conveyed in the Juniper Networks VSA, Downstream-Calculated-Qos-Rate (IANA 4874, 26-141). The receive speed is the recommended traffic value in bits per second used for upstream traffic for an ANCP interface, and is conveyed in the Juniper Networks VSA, Upstream-Calculated-Qos-Rate VSA (IANA 4874, 26-142).

To set the recommended shaping rates that are used as the default values for these VSAs in static configurations, include the `downstream-rate` and `upstream-rate` statements at the `[edit interfaces interface-name unit logical-unit-number advisory-options]` hierarchy level.

To configure the recommended rates on dynamically created VLAN interfaces, include the `upstream-rate` or `downstream-rate` statements at the `[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-unit advisory-options]` hierarchy level.

To configure the recommended rates on dynamically created ACL interface sets, include the `upstream-rate` or `downstream-rate` statements at the `[edit dynamic-profiles profile-name interface-set $junos-interface-set-name interfaces $junos-interface-ifd-name advisory-options]` hierarchy level.

ANCP Agent Keepalives for CoS

The ANCP agent sends a keepalive message to CoS at specific intervals. If CoS does not receive a keepalive in the expected time, it reverts the shaping rate changes it made in response to the ANCP agent. You can adjust how long CoS waits for a keepalive message by including the `maximum-helper-restart-time` statement at the `[edit protocols ancp]` hierarchy level. The interval between keepalive messages is automatically set to one-third the value of the maximum helper restart time. For example, if you set the maximum helper restart time to 120 seconds, then the ANCP agent sends keepalive messages every 40 seconds. In this example, if CoS does not receive a keepalive message within 120 seconds, then it reverts any policy changes derived from the ANCP agent.

Preservation of CoS Shaping Across ANCP Agent Restarts

When the ANCP agent stops due to a process or GRES, CoS enforces the ANCP downstream shaping-rates until the CoS keepalive timer expires. When the timer expires, CoS reverts to the CoS shaping-rate configured for the interfaces.

You configure the CoS keepalive timer by including the `maximum-helper-restart-time seconds` statement at the `[edit protocols ancp]` hierarchy level. It specifies how much time other daemons such as CoS wait for the ANCP agent to restart and is used to configure the CoS rate update keepalive timer.

The ANCP agent does not maintain TCP sessions from neighbors across the restart or GRES. When it restarts, it must re-establish sessions with neighbors and subscriber sessions before the timer expires. For all the re-established sessions, the ANCP agent updates CoS with the updated downstream shaping rates and provides DSL line attributes to the session database for AAA.

If CoS stops or restarts while ANCP is up, the ANCP agent retransmits all known subscriber downstream rates to CoS. Any existing adjusted shaping rates that have not been updated revert to the configured CoS shaping rates when the CoS restart timer expires.

SEE ALSO

[Specifying How Long Processes Wait for the ANCP Agent Restart to Complete](#) | 1648

Configuring the ANCP Agent to Report Traffic Rates to CoS

By default, the ANCP agent does not report the traffic rate on subscriber access lines to CoS. You must include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level to configure the ANCP agent to report downstream data rates received in ANCP Port Up messages to CoS for all subscribers in the network. This information enables CoS to subsequently shape the traffic on these access lines—but only if a shaping rate is configured in a CoS traffic-control profile for the access lines.

When an access node (DSLAM, ONT, ONU) calculates the data rate on the subscriber local loop, it ignores the additional headers on the subscriber access line that are associated with the overhead of the access mode (ATM or Ethernet). The unadjusted downstream data rate includes these headers in its calculation and reports a slightly higher value than that calculated by the access node. The ANCP agent also reports to CoS the traffic mode and the traffic rate overhead.



NOTE: The ANCP agent never reports upstream traffic rates to CoS.

You can also configure the ANCP agent to adjust the actual (net) downstream data rates for individual DSL types as follows:

- For frame-mode DSL types (G.fast, bonded G.fast, SDSL, bonded SDSL, VDSL, VDSL2, VDSL2 Annex Q, bonded VDSL2 Annex Q, and OTHER) and PON types (GPON, TWDM-PON, WDM-PON, XG-PON, and XGS-PON), you can configure an adjustment in the number of overhead bytes to account for encapsulation differences. You can also specify a percentage value that is applied to the actual, unadjusted data rate received in ANCP Port Up messages.
- For cell-mode DSL types (ADSL, ADSL2, and ADSL2+), you can configure only an adjustment in the number of overhead bytes for the traffic to account for encapsulation differences.

The ANCP agent adjusts the rate by the specified percentage. It adjusts the cell or frame overhead by adding or subtracting the specified number of bytes. By default the adjustment is 100 percent and 0 bytes, meaning that the agent does not adjust the net values before it reports them to CoS.

If CoS does not receive a keepalive message within the maximum helper restart time, it considers the ANCP agent to be down and immediately reverts to the configured values for any traffic shaping

updates that were modified as a result of traffic reports from the ANCP agent. The configured values are maintained until the ANCP agent comes back up and sends fresh traffic updates to CoS.

If the ANCP agent remains running but loses the connection to a neighbor, CoS does not revert to its configured values. In this case, CoS changes the shaping rate for the subscriber traffic only if the ANCP agent restores the connection to that neighbor and reports new traffic rates to CoS or if you remove the `qos-adjust` statement.



NOTE: Starting in Junos OS Release 17.4R1, the previously supported rate- and byte-adjustment statements at the `[edit protocols ancp]` and `[edit protocols ancp qos-adjust]` hierarchy levels are deprecated. They are replaced by the `access-line` statement and its many options at the `[edit system]` hierarchy level. The ANCP agent ignores the deprecated statements if they are present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing rate or byte adjustment configuration, you must reconfigure your adjustment with the `access-line` statement.

The exception to this change is that the `qos-adjust` statement remains supported, but no longer has any subordinate statements. Always configure the `qos-adjust` statement for normal ANCP operations. You may want to disable it for debugging purposes.



NOTE: Starting in Junos 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.



BEST PRACTICE: We recommend that you update your scripts to use the `ds/` statement when you upgrade to Junos OS Release 19.3R1 or higher releases. The redirect function will be supported for only a limited time.

To configure the ANCP agent to report traffic rates to CoS:

1. Enable rate reporting to CoS.

```
[edit protocols ancp]
user@host# set qos-adjust
```

2. (Optional) Specify the number of overhead bytes to add or subtract per cell or frame for one or more DSL line types.

- In Junos OS Release 17.3 or earlier:

```
[edit protocols ancp qos-adjust]
user@host# set adsl-bytes bytes
user@host# set adsl2-bytes bytes
user@host# set adsl2-plus-bytes bytes
user@host# set sds1-bytes bytes
user@host# set vds1-bytes bytes
user@host# set vds12-bytes bytes
user@host# set other-bytes bytes
```

- In Junos OS Release 17.4R1 or later:

```
[edit system access-line]
user@host# set adsl-overhead-bytes bytes
user@host# set adsl2-overhead-bytes bytes
user@host# set adsl2-plus-overhead-bytes bytes
user@host# set other-overhead-bytes bytes
user@host# set sds1-overhead-bytes bytes
user@host# set vds1-overhead-bytes bytes
user@host# set vds12-overhead-bytes bytes
```

- In Junos OS Release 18.4R1 or later, you can configure the following additional DSL line types:

```
[edit system access-line]
user@host# set gfast-bonded-overhead-bytes bytes
user@host# set gfast-overhead-bytes bytes
user@host# set sds1-bonded-overhead-bytes bytes
user@host# set vds12-annex-q-bonded-overhead-bytes bytes
user@host# set vds12-annex-q-overhead-bytes bytes
user@host# set vds12-bonded-overhead-bytes bytes
```

- In Junos OS Release 19.3R1 or later:

```
[edit system access-line dsl]
user@host# set adsl overhead-bytes bytes
user@host# set adsl2 overhead-bytes bytes
user@host# set adsl2-plus overhead-bytes bytes
user@host# set gfast overhead-bytes bytes
```

```

user@host# set gfast-bonded overhead-bytes bytes
user@host# set other overhead-bytes bytes
user@host# set sds1 overhead-bytes bytes
user@host# set sds1-bonded overhead-bytes bytes
user@host# set type tlv-value overhead-bytes bytes
user@host# set vds1 overhead-bytes bytes
user@host# set vds12 overhead-bytes bytes
user@host# set vds12-annex-q overhead-bytes bytes
user@host# set vds12-annex-q-bonded overhead-bytes bytes
user@host# set vds12-bonded overhead-bytes bytes

```

3. (Optional) In Junos OS Release 19.3R1 or later, specify the number of overhead bytes to add or subtract per cell or frame for one or more PON line types:

```

[edit system access-line pon]
user@host# set gpon overhead-bytes bytes
user@host# set other overhead-bytes bytes
user@host# set twdm-pon overhead-bytes bytes
user@host# set type tlv-value overhead-bytes bytes
user@host# set wdm-pon overhead-bytes bytes
user@host# set xg-pon overhead-bytes bytes
user@host# set xgs-pon overhead-bytes bytes

```

4. (Optional) Specify a percentage rate adjustment for one or more frame-mode DSL line types.

- In Junos OS Release 17.3 or earlier:

```

[edit protocols ancp qos-adjust]
user@host# set other-overhead-adjust percentage;
user@host# set sds1-overhead-adjust percentage
user@host# set vds1-overhead-adjust percentage
user@host# set vds12-overhead-adjust percentage;

```

- In Junos OS Release 17.4R1 or later:

```

[edit system access-line]
user@host# set other-overhead-adjust percentage
user@host# set sds1-overhead-adjust percentage
user@host# set vds1-overhead-adjust percentage
user@host# set vds12-overhead-adjust percentage

```

- In Junos OS Release 18.4R1 or later, you can configure the following additional DSL line types:

```
[edit system access-line]
user@host# set gfast-bonded-overhead-adjust percentage
user@host# set gfast-overhead-adjust percentage
user@host# set sds1-bonded-overhead-adjust percentage
user@host# set vds12-annex-q-overhead-adjust percentage
user@host# set vds12-annex-q-bonded-overhead-adjust percentage
user@host# set vds12-bonded-overhead-adjust percentage
```

- In Junos OS Release 19.3R1 or later:

```
[edit system access-line dsl]
user@host# set gfast overhead-adjust percentage
user@host# set gfast-bonded overhead-adjust percentage
user@host# set other overhead-adjust percentage
user@host# set sds1 overhead-adjust percentage
user@host# set sds1-bonded overhead-adjust percentage
user@host# set type tlv-value overhead-adjust percentage
user@host# set vds1 overhead-adjust percentage
user@host# set vds12 overhead-adjust percentage
user@host# set vds12-annex-q overhead-adjust percentage
user@host# set vds12-annex-q-bonded overhead-adjust percentage
user@host# set vds12-bonded overhead-adjust percentage
```

5. (Optional) In Junos OS Release 19.3R1 or later, specify a percentage rate adjustment for one or more PON line types:

```
[edit system access-line pon]
user@host# set gpon overhead-adjust percentage
user@host# set other overhead-adjust percentage
user@host# set twdm-pon overhead-adjust percentage
user@host# set type tlv-value overhead-adjust percentage
user@host# set wdm-pon overhead-adjust percentage
user@host# set xg-pon overhead-adjust percentage
user@host# set xgs-pon overhead-adjust percentage
```

SEE ALSO

[Shaping Rate Adjustments for Subscriber Local Loops Overview](#)

[Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops](#)

[Enabling Shaping-Rate Adjustments for Subscriber Local Loops](#)

[Disabling Shaping-Rate Adjustments for Subscriber Local Loops](#)

[Specifying How Long Processes Wait for the ANCP Agent Restart to Complete](#) | 1648

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces

When the access node sends information about the downstream and upstream calculated traffic rates for an interface, those values are used to shape the traffic sent to the interface so that it matches the subscriber local loop speed. You can specify recommended values that are used when the router does not receive this information from the access node. In this event, these recommended values are used as the default values for the following Juniper Networks VSAs:

- Downstream-Calculated-Qos-Rate (26-4874-141)—Conveys the transmit speed, which is the recommended traffic value in bits per second used for downstream traffic for an ANCP interface.
- Upstream-Calculated-Qos-Rate (26-4874-142)—Conveys the receive speed, which is the recommended traffic value in bits per second used for upstream traffic for an ANCP interface.

You can configure the recommended rates either on static VLAN and VLAN demux interfaces, or you can specify them in a dynamic profile for dynamic VLAN and VLAN demux interfaces or interface sets.

To configure recommended traffic shaping values for a static interface:

1. Set the rate in bits per second for downstream traffic for the interface.

```
[edit interfaces interface-name unit logical-unit-number advisory-options]
user@host# set downstream-rate rate
```

2. Set the rate in bits per second for upstream traffic for the interface.

```
[edit interfaces interface-name unit logical-unit-number advisory-options]
user@host# set upstream-rate rate
```

For example, to set the recommended downstream rate to 16 Mbps and the recommended upstream rate to 1 Mbps on VLAN demux interface demux0.10301:

```
[edit interfaces demux0 unit 10301 advisory-options]
user@host# set downstream-rate 16M
user@host# set upstream-rate 1M
```

To configure recommended traffic shaping values for a dynamic interface:

1. Set the rate in bits per second for downstream traffic in the dynamic profile.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-
interface-unit advisory-options]
user@host# set downstream-rate rate
```

2. Set the rate in bits per second for upstream traffic in the dynamic profile.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-
interface-unit advisory-options]
user@host# set upstream-rate rate
```

For example, to configure the dynamic profile ancp-dyn-vlan2 to set the recommended downstream rate to 10 Mbps and the recommended upstream rate to 1 Mbps on all interfaces in the dynamically created interface set:

```
[edit dynamic-profiles ancp-dyn-vlan2 interfaces $junos-interface-ifd-name unit $junos-interface-
unit advisory-options]
user@host# set downstream-rate 10M
user@host# set upstream-rate 1M
```

To configure recommended traffic shaping values for a dynamic interface set:

1. Set the rate in bits per second for downstream traffic in the dynamic profile.

```
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name
interface $junos-interface-ifd-name advisory-options]
user@host# set downstream-rate rate
```

2. Set the rate in bits per second for upstream traffic in the dynamic profile.

```
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name
interface $junos-interface-ifd-name advisory-options]
user@host# set upstream-rate rate
```

For example, to configure the dynamic profile `ancp-dyn-vlan1` to set the recommended downstream rate to 12 Mbps and the recommended upstream rate to 2 Mbps on all interfaces in the dynamically created interface set:

```
[edit dynamic-profiles ancp-dyn-vlan1 interfaces interface-set $junos-interface-set-name
interface $junos-interface-ifd-name advisory-options]
user@host# set downstream-rate 12M
user@host# set upstream-rate 2M
```

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates

The ANCP agent always reports both upstream and downstream rates to AAA. When a DSLAM calculates the data rate on the subscriber local loop, it ignores the additional headers on the DSL line that are associated with the overhead of the access mode (ATM or Ethernet). When the ANCP agent reports the net upstream data rate or the net downstream data rate, it includes the headers in its calculation and reports a slightly higher value than that calculated by the DSLAM; this is the unadjusted data rate.

The ANCP agent can optionally report adjusted data rates to AAA. Configure the agent to adjust the traffic rate to account for the header overhead by specifying an adjustment factor for one or more DSL line types. The adjustment factor is a percentage that is applied to the total downstream and upstream data rates reported by the ANCP agent. The adjustment factor applies globally for all subscribers of that DSL line type. By default, the ANCP agent applies an adjustment factor of 100 percent to all DSL lines, meaning that no adjustment is made. The ANCP agent simply passes on the DSL line rates that include the header information.



NOTE: These adjustment factors affect only the rates reported to AAA. The ANCP agent reports downstream data rates to CoS only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.



NOTE: Starting in Junos OS Release 17.4R1, the previously supported `qos-adjust-line-type` rate adjustment statements at the `[edit protocols ancp]` hierarchy level are deprecated. They are replaced by the `line-type-total-adjust` options for the access-line statement at the `[edit system]` hierarchy level. The ANCP agent ignores the deprecated statements if they are present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing rate adjustment configuration, you must reconfigure your adjustment with the access-line statement.

To apply a global adjustment factor for DSL subscriber lines to be reported to AAA:

- Specify the adjustment factor percentage for the desired subscriber line.
- In Junos OS Release 17.3 or earlier:

```
[edit protocols ancp]
user@host# set qos-adjust-adsl percentage
user@host# set qos-adjust-adsl2 percentage
user@host# set qos-adjust-adsl2-plus percentage
user@host# set qos-adjust-other percentage
user@host# set qos-adjust-sdsl percentage
user@host# set qos-adjust-vdsl percentage
user@host# set qos-adjust-vdsl2 percentage
```

- In Junos OS Release 17.4 or later:

```
[edit system access-line]
user@host# set adsl-total-adjust percentage
user@host# set adsl2-total-adjust percentage
user@host# set adsl2-plus-total-adjust percentage
user@host# set other-total-adjust percentage
user@host# set sdsl-total-adjust percentage
user@host# set vdsl-total-adjust percentage
user@host# set vdsl2-total-adjust percentage
```

- In Junos OS Release 18.4 or later, you can configure the following additional DSL line types:

```
[edit system access-line]
user@host# set gfast-bonded-total-adjust percentage
user@host# set gfast-total-adjust percentage
```

```

user@host# set sds1-bonded-total-adjust percentage
user@host# set vds12-annex-q-bonded-total-adjust percentage
user@host# set vds12-annex-q-total-adjust percentage
user@host# set vds12-bonded-total-adjust percentage

```

Starting in Junos 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.



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.

To apply a global adjustment factor for DSL subscriber lines to be reported to AAA in Junos OS Release 19.3R1 or later:

- Specify the adjustment factor percentage for the desired subscriber line.

```

[edit system access-line dsl]
user@host# set ads1 total-adjust percentage
user@host# set ads12 total-adjust percentage
user@host# set ads12-plus total-adjust percentage
user@host# set gfast total-adjust percentage
user@host# set gfast-bonded total-adjust percentage
user@host# set other total-adjust percentage
user@host# set sds1 total-adjust percentage
user@host# set sds1-bonded total-adjust percentage
user@host# set type tlv-value total-adjust percentage
user@host# set vds1 total-adjust percentage
user@host# set vds12 total-adjust percentage
user@host# set vds12-annex-q-bonded total-adjust percentage
user@host# set vds12-annex-q total-adjust percentage
user@host# set vds12-bonded total-adjust percentage

```

Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates

Starting in Junos Release 19.3R1, we support PON subscriber access line types in addition to the previously supported DSL line types.

The ANCP agent always reports both upstream and downstream rates to AAA. When an OLT or ONU access node calculates the data rate on the subscriber local loop, it ignores the additional headers on the

PON line that are associated with the overhead of the access mode (ATM or Ethernet). When the ANCP agent reports the net upstream data rate or the net downstream data rate, it includes the headers in its calculation and reports a slightly higher value than that calculated by the access node; this is the unadjusted data rate.

The ANCP agent can optionally report adjusted data rates to AAA. Configure the agent to adjust the traffic rate to account for the header overhead by specifying an adjustment factor for one or more PON line types. The adjustment factor is a percentage that is applied to the total downstream and upstream data rates reported by the ANCP agent. The adjustment factor applies globally for all subscribers of that PON line type. By default, the ANCP agent applies an adjustment factor of 100 percent to all PON lines, meaning that no adjustment is made. The ANCP agent simply passes on the DSL line rates that include the header information.

For PON line types, the adjustment is made to the total Layer 1 and encapsulation overhead in the following ANCP TLVs:

- ONT/ONU-Peak-Data-Rate-Downstream (0x94)
- ONT/ONU-Maximum-Data-Rate-Upstream (0x95)

The ANCP agent reports the adjusted value to the RADIUS server in Access-Request messages in the following Juniper Networks VSAs (vendor ID 4874):

- Downstream-Calculated-QoS-Rate (26-141)
- Upstream-Calculated-QoS-Rate (26-142)

The ANCP agent reports the adjusted value to an L2TP LNS in following AVPs:

- Tx Connect Speed (AVP 24 in ICCN message)
- Rx Connect Speed (AVP 38 in ICCN message)
- Connect Speed Update AVP 97 in CSUN message)

To apply a global adjustment factor for PON subscriber lines to be reported to AAA in Junos OS Release 19.3R1 or later:

- Specify the adjustment factor percentage for the desired subscriber line.

```
[edit system access-line pon]
user@host# set gpon total-adjust percentage
user@host# set other total-adjust percentage
user@host# set twdm-pon total-adjust percentage
user@host# set type tlv-value total-adjust percentage
user@host# set wdm-pon total-adjust percentage
```

```
user@host# set xg-pon total-adjust percentage
user@host# set xgs-pon total-adjust percentage
```

Verifying and Monitoring CoS for ANCP Subscribers

IN THIS SECTION

- Purpose | 1699
- Action | 1699

Purpose

View ANCP CoS state information:

Action

- To display summary information about the CoS state for all ANCP subscribers:

```
user@host> show ancp cos
```

- To display information about the CoS state for an ANCP subscriber specified by the ACI:

```
user@host> show ancp cos "port-2-11"
```

- To display the most recently updated CoS information:

```
user@host> show ancp cos last-update
```

- To display the CoS information that is pending (will be used to update the fields):

```
user@host> show ancp cos pending-update
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.3R1	Starting in Junos Release 19.3R1, the pre-existing adjustment options were renamed and placed in the new <code>dsl</code> stanza. The old DSL options are deprecated, but they redirect to the new location.
19.3R1	Starting in Junos Release 19.3R1, the pre-existing adjustment options were renamed and placed in the new <code>dsl</code> stanza. The old DSL options are deprecated, but they redirect to the new location.
19.3R1	Starting in Junos Release 19.3R1, we support PON subscriber access line types in addition to the previously supported DSL line types.
17.4R1	Starting in Junos OS Release 17.4R1, the ANCP agent can use access line information that it receives in the PPPoE intermediate agent (PPPoE-IA) tags.
17.4R1	Starting in Junos OS Release 17.4R1, the previously supported <code>rate-</code> and <code>byte-adjustment</code> statements at the <code>[edit protocols ancp]</code> and <code>[edit protocols ancp qos-adjust]</code> hierarchy levels are deprecated. They are replaced by the <code>access-line</code> statement and its many options at the <code>[edit system]</code> hierarchy level.
17.4R1	Starting in Junos OS Release 17.4R1, the previously supported <code>qos-adjust-line-type</code> rate adjustment statements at the <code>[edit protocols ancp]</code> hierarchy level are deprecated. They are replaced by the <code>line-type-total-adjust</code> options for the <code>access-line</code> statement at the <code>[edit system]</code> hierarchy level.

RELATED DOCUMENTATION

[ANCP Agent Neighbors and Operations | 1621](#)

[Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | 1713](#)

[Juniper Networks VSAs Supported by the AAA Service Framework | 19](#)

ANCP Agent and AAA

IN THIS SECTION

● [ANCP Agent Interactions with AAA | 1701](#)

- [ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | 1703](#)
- [Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | 1713](#)
- [Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | 1714](#)

ANCP Agent Interactions with AAA

The ANCP agent reports both unadjusted (net) data rates and adjusted data rates for subscriber traffic to AAA for RADIUS authentication and accounting of subscriber sessions. The adjusted data rate enables RADIUS to allocate the appropriate services (including *class of service*) to PPPoE sessions during authentication. The rate reports also enable RADIUS accounting to track the class of service actually provided for the PPPoE sessions, which in turn enables accurate billing for subscriber services.

The access nodes send ANCP DSL attributes in ANCP messages to the router, where they are stored in the shared database. AAA maps the ANCP DSL attributes to both the Juniper Networks DSL VSAs (used by RADIUS) and the DSL Forum VSA subattributes (also called the DSL Forum VSAs). RADIUS uses these attributes during authentication and accounting for PPPoE sessions on the subscriber access line. The attributes persist even when the ANCP session to a given node has ended, enabling RADIUS to later apply these attributes to new sessions on that subscriber access line. To remove the attributes, you must delete the interface or interface set for the access line from the ANCP agent configuration.

The RADIUS profile must be configured to include the `juniper-access-line-attributes` option, or AAA does not report the attributes to RADIUS. If the ANCP DSL attributes are unavailable, AAA maps the session's advisory upstream and downstream data rates (as configured on the session's underlying interface) to the Juniper Networks VSAs, Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141], respectively. AAA subsequently provides only these VSAs to RADIUS.

For successful authentication and accounting by RADIUS, AAA has to correlate PPPoE and DHCP IP demux sessions with their access lines and their associated DSL attributes. Some access nodes provide the ACI in PADI/PADR packets for the PPPoE sessions or in the DHCP discovery packets for DHCP IP demux sessions.

When the ACI is not provided in a 1:1 VLAN model with interface sets, you must associate the underlying interface for the sessions with the identifier and the interface set. If you do not configure this association, then only the advisory traffic rates are provided to RADIUS. This configuration has no effect when the identifier is provided by the access node.

For the N:1 VLAN model with interface sets, the access node must provide the ACI. If you configure the underlying interface for this model when the access node does not provide the identifier, the subscriber sessions can be incorrectly correlated with access lines.

AAA reports values to RADIUS for the Juniper Networks VSAs 26-141 and 26-142 according to the following scheme:

1. When the PPPoE or DHCP IP demux subscriber session can be correlated with an access line, then the ANCP agent adjusts the downstream and upstream traffic rates reported by the access node according to the ANCP agent CoS configuration. The agent then maps the adjusted rates to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141].
2. If the session cannot be correlated with an access line, but the PPPoE or DHCP discovery packet includes the DSL Forum VSA and the Access-Loop-Encapsulation subattribute includes a value for the AAL5 data link, then the ANCP agent adjusts the Actual-Data-Rate-Downstream and Actual-Data-Rate-Upstream subattributes to account for the ATM 48/53 cell tax. The adjusted rates mapped to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141].
3. If neither of the preceding sets of conditions is satisfied, then the ANCP agent simply maps the recommended downstream and upstream data rates to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141]. The recommended rates are either configured statically for the VLAN or VLAN demux interfaces or are in the dynamic profile that creates the interfaces.

To map an ACI to a static VLAN demux interface, include the `access-identifier identifier` statement at the `[edit protocols ancp interfaces demux0.logical-unit-number]` hierarchy level.

To configure advisory upstream and downstream data rates on a static VLAN demux interface, include the `upstream-rate rate` or `downstream-rate rate` statements at the `[edit interfaces demux0 unit logical-unit-number]` hierarchy level.

To configure an underlying interface for the PPPoE sessions in an interface set, include the `underlying-interface interface-name` statement at the `[edit protocols ancp interfaces interface-set interface-set-name]` hierarchy level.

When an ACI, and therefore a subscriber access line, has been mapped to an interface or interface set, the ACI can be re-mapped to a different interface or set. When this happens, traffic shaping is adjusted accordingly for the interfaces or interface sets involved. This capability is useful for the Business Services model, where a PPPoE session that is initially classified as a residential household can be reclassified as a business subscriber during RADIUS authentication by using a Junos OS ICE AAA framework Op-Script application.

In the Business Services Model, the PPPoE session initially represents a residential household until RADIUS authentication and authorization takes place. The ANCP agent dynamically maps the household's access line to the appropriate subscriber interface and applies CoS traffic shaping to the interface. During authentication and authorization, the Op-Script application may classify the PPPoE session as a business subscriber rather than a residential subscriber. If this occurs, the application creates multiple static VLANs and groups them into an interface set. Based on the ANCP agent configuration, the application then statically maps the subscriber's access line to this static interface set. This interface set can include only static interfaces.

The ANCP agent reverts CoS traffic shaping from the interface previously used by the subscriber and instead applies the shaping to the interface set. This reversion means that the CoS process applies to the interface the next shaping rate in its adjustment control profile.

ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes

Some broadband access line information is not supported by standard RADIUS attributes. The DSL Forum defined RADIUS vendor-specific attributes for DSL access lines in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*. The VSAs include information about the access lines, the subscribers using the lines, and data rates on the lines.

The DSL Forum changed its name to the Broadband Forum and defined new RADIUS VSAs for G.fast (DSL) and PON access technologies. Some of the VSAs previously used only for DSL networks are also used for PON networks. All these VSAs, regardless of access technology, are referred to as DSL Forum VSAs because they are subattributes contained in the DSL Forum VSA.

An ANCP access node can provide this information to the router in the following ways:

- In ANCP messages that carry ANCP access line TLVs (Type-Length-Value attributes)
- In a PPPoE PADI message during PPPoE subscriber discovery

The original ANCP DSL TLVs are defined in RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*. RFC 6320 Draft Extension, *Access Extensions for the Access Node Control Protocol*, adds new TLVs for the DSL G.fast and PON VSAs. The ANCP access line TLVs map to both DSL Forum VSAs (IANA vendor ID 3561) and Juniper Networks (IANA vendor ID 4874) access line VSAs.

When the router receives ANCP TLVs from the access node, it does not parse or manipulate the information. Instead it simply passes the access line and traffic information to the RADIUS server in the corresponding RADIUS VSAs mapped from the TLVs. A RADIUS authentication or accounting message can contain any combination of the DSL Forum VSAs and the Juniper Networks VSAs. You can configure the RADIUS access profile to exclude one or more individual attributes, or all DSL Forum VSAs, from being included in RADIUS messages.

The DSL Forum VSAs received by the router during PPPoE and DHCP client discovery are not updated after discovery, whereas the equivalent ANCP attributes are updated whenever there is a change to the access line.

[Table 71 on page 1704](#) shows the relationship between the ANCP TLVs, Juniper Networks VSAs, and DSL Forum VSAs.

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0x01 Access-Loop-Circuit-ID	26-4874-110 Acc-Loop-Cir-Id	26-3561-1 Agent-Circuit-Id
0x02 Access-Loop-Remote-ID	26-4874-182 Acc-Loop-Remote-Id	26-3561-2 Agent-Remote-Id
0x03 Access-Aggregation-Circuit-ID-ASCII	26-4874-112 Acc-Aggr-Cir-Id-Asc	26-3561-3 Access-Aggregation-Circuit-ID-ASCII
0x06 Access-Aggregation-Circuit-ID-Binary	26-4874-111 Acc-Aggr-Cir-Id-Bin	26-3561-6 Access-Aggregation-Circuit-ID-Binary
0x81 Actual-Net-Data-Rate-Upstream	<ul style="list-style-type: none"> 26-4874-92 L2C-Up-Stream-Data—Unadjusted rate 26-4874-113 Act-Data-Rate-Up—Unadjusted rate 26-4874-142 Upstream-Calculated-Qos-Rate—Rate as adjusted by ANCP 	26-3561-129 Actual-Data-Rate-Upstream

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs (Continued)

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0x82 Actual-Net-Data-Rate-Downstream	<ul style="list-style-type: none"> 26-4874-93 L2C-Down-Stream-Data—Unadjusted rate 26-4874-114 Act-Data-Rate-Dn—Unadjusted rate 26-4874-141 Downstream-Calculated-Qos-Rate—Rate as adjusted by ANCP 	26-3561-130 Actual-Data-Rate-Downstream
0x83 Minimum-Net-Data-Rate-Upstream	26-4874-115 Min-Data-Rate-Up	26-3561-131 Minimum-Data-Rate-Upstream
0x84 Minimum-Net-Data-Rate-Downstream	26-4874-116 Min-Data-Rate-Dn	26-3561-132 Minimum-Data-Rate-Downstream
0x85 Attainable-Net-Data-Rate-Upstream	26-4874-117 Att-Data-Rate-Up	26-3561-133 Attainable-Data-Rate-Upstream
0x86 Attainable-Net-Data-Rate-Downstream	26-4874-118 Att-Data-Rate-Dn	26-3561-134 Attainable-Data-Rate-Downstream
0x87 Maximum-Net-Data-Rate-Upstream	26-4874-119 Max-Data-Rate-Up	26-3561-135 Maximum-Data-Rate-Upstream

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs (Continued)

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0x88 Maximum-Net-Data-Rate-Downstream	26-4874-120 Max-Data-Rate-Dn	26-3561-136 Maximum-Data-Rate-Downstream
0x89 Minimum-Net-Low-Power-Data-Rate-Upstream	26-4874-121 Min-LP-Data-Rate-Up	26-3561-137 Minimum-Data-Rate-Upstream-Low-Power
0x8A Minimum-Net-Low-Power-Data-Rate-Downstream	26-4874-122 Min-LP-Data-Rate-Dn	26-3561-138 Minimum-Data-Rate-Downstream-Low-Power
0x8B Maximum-Interleaving-Delay-Upstream	26-4874-123 Max-Interlv-Delay-Up	26-3561-139 Maximum-Interleaving-Delay-Upstream
0x8C Actual-Interleaving-Delay-Upstream	26-4874-124 Act-Interlv-Delay-Up	26-3561-140 Actual-Interleaving-Delay-Upstream
0x8D Maximum-Interleaving-Delay-Downstream	26-4874-125 Max-Interlv-Delay-Dn	26-3561-141 Maximum-Interleaving-Delay-Downstream
0x8E Actual-Interleaving-Delay-Downstream	26-4874-126 Act-Interlv-Delay-Dn	26-3561-142 Actual-Interleaving-Delay-Downstream

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs (Continued)

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0x8F DSL-Line-State	26-4874-127 DSL-Line-State	n/a
0x90 Access-Loop-Encapsulation	26-4874-183 Acc-Loop-Encap	26-3561-144 Access-Loop-Encapsulation
0x91 DSL-Type	26-4874-128 DSL-Type	26-3561-145 DSL-Type
0x92 PON-Access-Type	26-4874-219 PON-Access-Type	26-3561-146 PON-Access-Type
0x93 ONT/ONU-Average-Data-Rate-Downstream	26-4874-220 ONT/ONU-Average-Data-Rate-Downstream	26-3561-147 ONT/ONU-Average-Data-Rate-Downstream
0x94 ONT/ONU-Peak-Data-Rate-Downstream	26-4874-221 ONT/ONU-Peak-Data-Rate-Downstream	26-3561-148 ONT/ONU-Peak-Data-Rate-Downstream
0x95 ONT/ONU-Maximum-Data-Rate-Upstream	26-4874-222 ONT/ONU-Maximum-Data-Rate-Upstream	26-3561-149 ONT/ONU-Maximum-Data-Rate-Upstream
0x96 ONT/ONU-Assured-Data-Rate-Upstream	26-4874-223 ONT/ONU-Assured-Data-Rate-Upstream	26-3561-150 ONT/ONU-Assured-Data-Rate-Upstream

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs (Continued)

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0x97 PON-Tree-Maximum-Data-Rate-Upstream	26-4874-224 PON-Tree-Maximum-Data-Rate-Upstream	26-3561-151 PON-Tree-Maximum-Data-Rate-Upstream
0x98 PON-Tree-Maximum-Data-Rate-Downstream	26-4874-225 PON-Tree-Maximum-Data-Rate-Downstream	26-3561-152 PON-Tree-Maximum-Data-Rate-Downstream
0x9B Expected Throughput	26-4874-226 Expected-Throughput-Upstream	26-3561-155 Expected-Throughput-Upstream
0x9C Expected Throughput at L2	26-4874-227 Expected-Throughput-Downstream	26-3561-156 Expected-Throughput-Downstream
0x9D Attainable Expected Throughput	26-4874-228 Attainable-Expected-Throughput-Upstream	26-3561-157 Attainable-Expected-Throughput-Upstream
0x9E Attainable Expected Throughput at L2	26-4874-229 Attainable-Expected-Throughput-Downstream	26-3561-158 Attainable-Expected-Throughput-Downstream
0x9F Gamma data rate upstream	26-4874-230 Gamma-Data-Rate-Upstream	26-3561-159 Gamma-Data-Rate-Upstream
0xA0 Gamma data rate downstream	26-4874-231 Gamma-Data-Rate-Downstream	26-3561-160 Gamma-Data-Rate-Downstream

Table 71: Mapping Access Line Attributes: ANCP TLVs to Juniper VSAs to DSL Forum VSAs (Continued)

ANCP TLV Number and Name	Juniper Networks VSA Number and Name	DSL Forum VSA Number and Name
0xA1 Attainable Gamma data rate upstream	26-4874-232 Attainable-Gamma-Data-Rate- Upstream	26-3561-161 Attainable-Gamma-Data-Rate- Upstream
0xA2 Attainable Gamma data rate downstream	26-4874-233 Attainable-Gamma-Data-Rate- Downstream	26-3561-162 Attainable-Gamma-Data-Rate- Downstream

[Table 72 on page 1709](#) lists the ANCP TLVs and indicates with a checkmark whether the TLV is used for DSL or PON subscriber access lines.

Table 72: DSL and PON Support for ANCP TLVs

ANCP TLV Number and Name	Used for DSL Access	Used for PON Access
0x01 Access-Loop-Circuit-ID	✓	✓
0x02 Access-Loop-Remote-ID	✓	✓
0x03 Access-Aggregation-Circuit-ID-ASCII	✓	✓
0x06 Access-Aggregation-Circuit-ID-Binary	✓	✓

Table 72: DSL and PON Support for ANCP TLVs (Continued)

ANCP TLV Number and Name	Used for DSL Access	Used for PON Access
0x81 Actual-Net-Data-Rate-Upstream	✓	–
0x82 Actual-Net-Data-Rate-Downstream	✓	–
0x83 Minimum-Net-Data-Rate-Upstream	✓	–
0x84 Minimum-Net-Data-Rate-Downstream	✓	–
0x85 Attainable-Net-Data-Rate-Upstream	✓	–
0x86 Attainable-Net-Data-Rate-Downstream	✓	–
0x87 Maximum-Net-Data-Rate-Upstream	✓	–
0x88 Maximum-Net-Data-Rate-Downstream	✓	–
0x89 Minimum-Net-Low-Power-Data-Rate-Upstream	✓	–

Table 72: DSL and PON Support for ANCP TLVs (Continued)

ANCP TLV Number and Name	Used for DSL Access	Used for PON Access
0x8A Minimum-Net-Low-Power-Data-Rate-Downstream	✓	–
0x8B Maximum-Interleaving-Delay-Upstream	✓	–
0x8C Actual-Interleaving-Delay-Upstream	✓	–
0x8D Maximum-Interleaving-Delay-Downstream	✓	–
0x8E Actual-Interleaving-Delay-Downstream	✓	–
0x8F DSL-Line-State	✓	–
0x90 Access-Loop-Encapsulation	✓	–
0x91 DSL-Type	✓	–
0x92 PON-Access-Type	–	

Table 72: DSL and PON Support for ANCP TLVs (Continued)

ANCP TLV Number and Name	Used for DSL Access	Used for PON Access
0x93 ONT/ONU-Average-Data-Rate-Downstream	-	✓
0x94 ONT/ONU-Peak-Data-Rate-Downstream	-	✓
0x95 ONT/ONU-Maximum-Data-Rate-Upstream	-	✓
0x96 ONT/ONU-Assured-Data-Rate-Upstream	-	✓
0x97 PON-Tree-Maximum-Data-Rate-Upstream	-	✓
0x98 PON-Tree-Maximum-Data-Rate-Downstream	-	✓
0x9B Expected Throughput	✓	-
0x9C Expected Throughput at L2	✓	-
0x9D Attainable Expected Throughput	✓	-

Table 72: DSL and PON Support for ANCP TLVs (Continued)

ANCP TLV Number and Name	Used for DSL Access	Used for PON Access
0x9E Attainable Expected Throughput at L2	✓	–
0x9F Gamma data rate upstream	✓	–
0xA0 Gamma data rate downstream	✓	–
0xA1 Attainable Gamma data rate upstream	✓	–
0xA2 Attainable Gamma data rate downstream	✓	–

Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages

You can include the `juniper-access-line-attributes` statement to configure AAA to add the set of Juniper Networks access line VSAs to the RADIUS authentication and accounting request messages for subscribers. By default, these VSAs are not added to any RADIUS message. See ["ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes" on page 1703](#) for a table of the Juniper Networks DSL VSAs.

After you have configured the inclusion of the Juniper Networks access line VSAs, you can subsequently exclude one or more of the VSAs from being transmitted. To do so, include the `exclude` statement at the `[edit access profile profile-name radius attributes]` hierarchy level, and specify which VSAs to exclude.

In contrast to the Juniper Networks access line VSAs (vendor ID 4874), the DSL Forum VSA (vendor ID 3561) is added to all RADIUS messages by default. The DSL Forum VSA conveys individual DSL Forum attributes. See ["DSL Forum Vendor-Specific Attributes" on page 626](#) for a table of these VSAs. You can use the `exclude` statement at the `[edit access profile profile-name radius attributes]` hierarchy level to prevent this VSA from being included in any RADIUS message.

To add the Juniper Networks access line VSAs to RADIUS messages:

- Configure the inclusion trigger.

```
[edit access profile profile-name radius options]
user@host# set juniper-access-line-attributes
```

To exclude specific Juniper Networks DSL VSAs from RADIUS messages:

- Configure the exclusion trigger.

```
[edit access profile profile-name radius attributes]
user@host# set exclude vsa-option
```

For example, to exclude the interleaving delay VSAs, configure the following statements:

```
[edit access profile profile-name radius attributes]
user@host# set exclude max-interlv-delay-dn
user@host# set excludemax-interlv-delay-up
```

To exclude the DSL Forum (RFC 4679) VSA from RADIUS messages:

- Configure the exclusion trigger.

```
[edit access profile profile-name radius attributes]
user@host# set exclude dsl-forum-attributes
```

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications

When an ANCP neighbor reports a change in the upstream traffic rate or downstream traffic rate of an access line, the ANCP agent immediately passes the information to AAA. By default, AAA does not pass this information on to the RADIUS server until the next accounting update. However, you can configure AAA to report the rate change immediately.

When you include the `ancp-speed-change-immediate-update` statement in the subscriber session access profile, receipt of the notification from the ANCP agent triggers AAA to send an interim update Accounting-Request message to the RADIUS server for the PPPoE and DHCP IP demux subscribers associated with that access line. The interim update request includes the new access line parameters and the adjusted upstream and downstream traffic rates.

To configure AAA to immediately send rate change information from the ANCP agent to the RADIUS server with interim accounting updates:

- Specify the immediate update.

```
[edit access profile profile-name accounting]  
user@host# set ancp-speed-change-immediate-update
```

SEE ALSO

[Configuring Per-Subscriber Session Accounting | 744](#)

RELATED DOCUMENTATION

[ANCP Agent Neighbors and Operations | 1621](#)

[Juniper Networks VSAs Supported by the AAA Service Framework | 19](#)

ANCP Monitoring and Management

IN THIS SECTION

- [Triggering ANCP OAM to Test the Local Loop | 1715](#)
- [Verifying and Monitoring ANCP Neighbors | 1717](#)
- [Clearing ANCP Neighbors | 1718](#)
- [Verifying and Monitoring ANCP Subscribers | 1719](#)
- [Clearing ANCP Subscribers | 1720](#)
- [Clearing and Verifying ANCP Statistics | 1721](#)

Triggering ANCP OAM to Test the Local Loop

You can trigger ANCP OAM to perform a loopback test on the local loop between the access node and the CPE to help isolate simple faults. On an ATM-based local loop, the ANCP operation triggers the access node to generate ATM (F4/F5) loopback cells on the local loop. On an Ethernet-based local loop,

the ANCP operation triggers the access node to generate an Ethernet loopback message on the local loop. When the test completes, the access node sends a message to the router with the results.

Issue the `request ancp oam neighbor` command from CLI operational mode to initiate testing of a local loop identified by the IP address or system name of the ANCP neighbor and the ACI for a subscriber on that access node.

Issue the `request ancp oam interface` command from CLI operational mode to initiate testing of a local loop identified by the ANCP interface or interface set associated with a subscriber and the ACI for a subscriber on that access node.

With both commands, you can also specify how many times the test must be run and how long the router waits for a response to the OAM request.

To initiate ANCP local loop testing:

- Identify the loop by the subscriber identifier and the neighbor's IP address; optionally specify how many times the test runs and the timeout period.

```
user@host> request ancp oam neighbor ip-address 192.168.32.5 subscriber "dslam port-2-10"
count 5 timeout 600
```

- Identify the loop by the subscriber identifier and the neighbor's system name; optionally specify how many times the test runs and the timeout period.

```
user@host> request ancp oam neighbor system-name 00:00:5E:00:53:ba subscriber "dslam
port-2-10" count 10 timeout 600
```

- Identify the loop by the subscriber identifier and the interface associated with the subscriber; optionally specify how many times the test runs and the timeout period.

```
user@host> request ancp oam interface ge-1/0/2.12 identifier-string timeout 15
```

- Identify the loop by the subscriber identifier and the set of interfaces associated with the subscriber; optionally specify how many times the test runs and the timeout period.

```
user@host> request ancp oam interface interface-set vlan5 identifier-string count 3
```

Verifying and Monitoring ANCP Neighbors

IN THIS SECTION

- Purpose | 1717
- Action | 1717

Purpose

View ANCP neighbor information:

Action

- To display summary information about all ANCP neighbors:

```
user@host> show ancp neighbor
```

- To display information about a specific ANCP neighbor, add the IP address or MAC address to the command:

```
user@host> show ancp neighbor ip-address 203.0.113.64
user@host> show ancp neighbor system-name 00:00:5E:00:53:ba
```

- To display detailed information, add **detail** to the command:

```
user@host> show ancp neighbor detail
```

```
user@host> show ancp neighbor system-name 00:00:5E:00:53:ba detail
```

- To display a count of ANCP neighbors in various states and the total number of neighbors, or a count of DSL lines in various states for all subscribers for a particular neighbor:

```
user@host> show ancp summary neighbor
user@host> show ancp summary neighbor 203.0.113.64
```

- To display total and state-wise counts of both ANCP neighbors and subscribers:

```
user@host> show ancp summary
```

Clearing ANCP Neighbors

IN THIS SECTION

- Purpose | 1718
- Action | 1718

Purpose

Clear ANCP neighbor information.

Action

- To clear connections with all ANCP neighbors:

```
user@host> clear ancp neighbor
```

- To clear the connection with a specific ANCP neighbor, add the IP address or MAC address to the command:

```
user@host> clear ancp neighbor ip-address 203.0.113.64
```

```
user@host> clear ancp neighbor system-name 00:00:5E:00:53:ba
```

- To verify that the connection has been cleared:

```
user@host> show ancp neighbor
```

```
user@host> show ancp neighbor 203.0.113.64
```

```
user@host> show ancp neighbor 00:00:5E:00:53:ba
```

Verifying and Monitoring ANCP Subscribers

IN THIS SECTION

- [Purpose | 1719](#)
- [Action | 1719](#)

Purpose

View ANCP subscriber (local access loop) information:

Action

- To display summary information about all ANCP subscribers:

```
user@host> show ancp subscriber
```

- To display information about all ANCP subscribers connected through a particular ANCP neighbor:

```
user@host> show ancp subscriber neighbor 203.0.113.64
```

- To display information about an ANCP subscriber specified by the ACI:

```
user@host> show ancp subscriber "port-2-11"
```

- To display detailed information, add **detail** to the command:

```
user@host> show ancp subscriber detail
```

```
user@host> show ancp subscriber neighbor 203.0.113.64 detail
```

- To display a count of subscribers in various states and the total number of subscribers:

```
user@host> show ancp summary subscriber
```

- To display total and state-wise counts of both ANCP neighbors and subscribers:

```
user@host> show ancp summary
```

Clearing ANCP Subscribers

IN THIS SECTION

- [Purpose | 1720](#)
- [Action | 1720](#)

Purpose

Clear ANCP subscriber information.

Action

- To clear connections with all ANCP subscribers that are not mapped:

```
user@host> clear ancp subscriber
```

- To clear connections with all ANCP subscribers that are mapped:

```
user@host> clear ancp neighbor
```

- To clear the connection with an ANCP subscriber identified by a particular ACI on all neighbors, add the identifier to the command:

```
user@host> clear ancp subscriber identifier port-2-10
```

- To clear the connection with an ANCP subscriber identified by a particular ACI on a specific neighbor, add the identifier and either the IP address or MAC address to the command:

```
user@host> clear ancp subscriber identifier port-2-10 ip-address 203.0.113.64
```

```
user@host> clear ancp subscriber identifier port-2-10 system-name 00:00:5E:00:53:ba
```

- To verify that the connection has been cleared:

```
user@host> show ancp subscriber
```

Clearing and Verifying ANCP Statistics

IN THIS SECTION

● Purpose | 1721

● Action | 1722

Purpose

Clear ANCP statistics.

Action

- To clear all ANCP statistics:

```
user@host> clear ancp statistics
```

- To clear statistics for a specific ANCP neighbor, add the IP address or MAC address to the command:

```
user@host> clear ancp statistics ip-address 203.0.113.64
```

```
user@host> clear ancp statistics system-name 00:00:5E:00:53:ba
```

- To verify that the statistics have been cleared:

```
user@host> show ancp statistics
```

RELATED DOCUMENTATION

| [ANCP Agent Neighbors and Operations](#) | 1621

Tracing ANCP Events for Troubleshooting

IN THIS SECTION

- [Configuring the ANCP Trace Log Filename](#) | 1723
- [Configuring the Number and Size of ANCP Log Files](#) | 1723
- [Configuring Access to the ANCP Log File](#) | 1724
- [Configuring a Regular Expression for ANCP Messages to Be Logged](#) | 1724
- [Configuring the ANCP Tracing Flags](#) | 1725
- [Configuring the Severity Level to Filter Which ANCP Messages Are Logged](#) | 1725

The Junos OS trace feature tracks ANCP agent 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 `ancpd`. 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 ANCP agent operations:

Configuring the ANCP Trace Log Filename

By default, the name of the file that records trace output for ANCP is `ancpd`. You can specify a different name with the `file` option.

To configure the filename for ANCP tracing operations:

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

```
[edit protocols ancp traceoptions]
user@host# set file ancp_1
```

Configuring the Number and Size of ANCP 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 ancp traceoptions]
user@host# set file ancp_1 _logfile_1 files 20 size 2097152
```

Configuring Access to the ANCP 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 ancp traceoptions]
user@host# set file ancp_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 ancp traceoptions]
user@host# set file ancp_1 _logfile_1 no-world-readable
```

Configuring a Regular Expression for ANCP 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 protocols ancp traceoptions]  
user@host# set file ancp_1 _logfile_1 match regex
```

Configuring the ANCP 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 ancp traceoptions]  
user@host# set flag restart
```

Configuring the Severity Level to Filter Which ANCP 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. 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` or `verbose`. 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 ancp traceoptions]  
user@host# set level severity
```

RELATED DOCUMENTATION

| [ANCP Agent Neighbors and Operations](#) | 1621

10

PART

Diameter Base Protocol and its Applications

- [Diameter Base Protocol and its Applications | 1727](#)
-

Diameter Base Protocol and its Applications

IN THIS CHAPTER

- Diameter Base Protocol | 1727
- Gx-Plus for Provisioning Subscribers | 1781
- 3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1799
- NASREQ for Authentication and Authorization | 1852
- JSRC for Subscriber Provisioning and Accounting | 1856
- JSRC and Subscribers on Static Interfaces | 1872
- Monitoring and Management Diameter Information | 1888
- Tracing Diameter Base Protocol Events for Troubleshooting | 1895
- Troubleshooting Diameter Networks | 1899
- Monitoring and Managing Static Subscriber Information | 1901
- Tracing Static Subscriber Events for Troubleshooting | 1903

Diameter Base Protocol

IN THIS SECTION

- Diameter Base Protocol Overview | 1728
- Messages Used by Diameter Applications | 1731
- Diameter AVPs and Diameter Applications | 1739
- Configuring Diameter | 1762
- Configuring the Origin Attributes of the Diameter Instance | 1763
- Configuring Diameter Peers | 1763
- Configuring the Diameter Transport | 1765
- Configuring Diameter Network Elements | 1766

- [Example: Configure S6a Application | 1768](#)

Diameter Base Protocol Overview

IN THIS SECTION

- [Benefits of Using Diameter | 1730](#)

The Diameter protocol is defined in *RFC 3588, Diameter Base Protocol*, and provides an alternative to RADIUS that is more flexible and extensible. The Diameter base protocol provides basic services to one or more applications (also called functions) that runs in a different Diameter instance. The individual application provides the extended AAA functionality. Applications that use Diameter include Gx-Plus, JSRC, NASREQ, PTSP, and S6a. Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Diameter peers communicate over a reliable TCP transport layer connection by exchanging Diameter messages that convey status, requests, and acknowledgments by means of standard Diameter AVPs and application-specific AVPs. The Diameter transport layer configuration is based on Diameter network elements (DNEs); multiple DNEs per Diameter instance are supported. Currently only the predefined *master* Diameter instance is supported, but you can configure alternative values for many of the master Diameter instance values.

Each DNE consists of a prioritized list of peers and a set of routes that define how traffic is forwarded. Each route associates a destination with a function (application), a function partition, and a metric. When an application sends a message to a routed destination, all routes within the Diameter protocol instance are examined for a match. When the best route to the destination has been selected, the message is forwarded by means of the DNE that includes that route.

Multiple routes to the same destination can exist within a given DNE and in different DNEs. In the case of multiple routes that match a request for forwarding, the best route is selected as follows:

1. The route with the lowest metric is selected.
2. In the event of a tie, the route with the highest specification score is selected.
3. In the event of another tie, then the names of the DNEs are compared in lexicographical order. The route in the DNE with the lowest value is selected. For example, dne-austin has a lower value than dne-boston.

4. If the routes are tied within the same DNE, then the route names are compared in lexicographical order. The route with the lowest value is selected.

The specification score of a route is 0 by default. Points are added to the score as follows:

- If the destination realm matches the request, add 1.
- If the destination host matches the request, add 2.
- If the function matches the request, add 3.
- If the function partition matches the request, add 4.

Multiple routes to the same destination can exist within a given DNE and in different DNEs. In the case of multiple routes that match a request for forwarding, Diameter selects the best route as follows:

1. Diameter compares the metric of the routes and selects the route with the lowest metric.
2. If multiple routes have the same lowest metric, then Diameter selects the most-qualified route. Diameter evaluates multiple attributes of the route to determine a score that reflects how specifically each route matches the request. By default, the score of a route is 0. Points are added to the score as follows:
 - If the destination realm matches the request, add 1.
 - If the destination host matches the request, add 2.
 - If the function matches the request, add 3.
 - If the function partition matches the request, add 4.
3. If multiple routes are equally qualified, then Diameter compares the names of the DNEs in lexicographical order and selects the route in the DNE that has the lowest value. For example, dne-austin has a lower value than dne-boston.
4. If the routes are tied within the same DNE, then Diameter compares the route names in lexicographical order and selects the route with the lowest value.

When the state of any DNE changes, the route lookup for all destinations is reevaluated. All outstanding messages to routed destinations are rerouted as needed, or discarded.

To configure a Diameter network element, include the `network-element` statement at the `[edit diameter]` hierarchy level, then include the `route` statement at the `[edit diameter network-element element-name forwarding]` hierarchy level.

To configure a route for the DNE, include the `destination` (optional), `function` (optional), and `metric` statements at the `[edit diameter network-element element-name forwarding route dne-route-name]` hierarchy level.

Specify the Diameter peers associated with the DNE by including one or more `peer` statements at the `[edit diameter network-element element-name]` hierarchy level.

Set the priority for each peer with the `priority` statement at the `[edit diameter network-element element-name peer peer-name]` hierarchy level.

Diameter requires you to configure information about the origin node; this is the endpoint node that originates Diameter for the Diameter instance. Include the `host` and `realm` statements at the `[edit diameter]` hierarchy level to configure the Diameter origin.

You can optionally configure one or more *transports* to specify the source (local) address of the transport layer connection. To configure a Diameter transport, include the `transport` statement at the `[edit diameter]` hierarchy level. Then include the `address` statement at the `[edit diameter transport transport-name]` hierarchy level.

You can optionally specify a logical system and routing instance for the connection by including the `logical-system` and `routing-instance` statements at the `[edit diameter transport transport-name]` hierarchy level. By default, Diameter uses the default logical system and default routing instance (using the main inet.0 routing table). The logical system and routing instance for the transport connection must match that for the peer, or a configuration error is reported.

Each Diameter peer is specified by a name. Peer attributes include address and the destination TCP port used by active connections to this peer. To configure a Diameter peer, include the `peer` statement at the `[edit diameter]` hierarchy level, and then include the `address` and `connect-actively` statements at the `[edit diameter peer peer-name]` hierarchy level.

To configure the active connection, include the `port` and `transport` statements at the `[edit diameter peer peer-name connect-actively]` hierarchy level. The assigned transport identifies the transport layer source address used to establish active connections to the peers. `transport` statements.

Benefits of Using Diameter

- Diameter enables a lower load on the network and servers by reporting usage information at a much lower frequency compared to RADIUS. RADIUS involves periodic updates independent of usage changes. Diameter applications such as Gx enable you to set thresholds with correlating pushes of usage statistics from the router to the PCRF. The PCRF can then make appropriate adjustments to services and costs.
- Wireless services and charging are typically performed with Diameter applications, but wireline services have generally used a RADIUS-based infrastructure. Customers with both wireline and wireless offerings can reduce the complexity and cost of maintaining separate infrastructures by migrating their wireline operations to their existing Diameter-based wireless infrastructure.
- Applications that run over Diameter tend to be stateful (some may be either, such as NASREQ), whereas RADIUS is not stateful.

- Multiple application protocols can run over Diameter, such as NASREQ, Gx, Gy, JSRC, and S6a.
- Larger attribute space than RADIUS, which enables a greater number of standard and vendor-specific attributes (AVPs) than RADIUS. Diameter also supports the RADIUS standard attributes, reserving AVPs 1 through 255 for them.

Messages Used by Diameter Applications

Junos OS supports the following Diameter applications:

- JSRC—A Juniper Networks Diameter application registered with the IANA (<http://www.iana.org>) as Juniper Policy-Control-JSRC, with an ID of 16777244. It communicates with the SAE (remote SRC peer).
- PTSP—A Juniper Networks Diameter application registered with the IANA (<http://www.iana.org>) as Juniper JGx, with an ID of 16777273. It communicates with the SAE (remote SRC peer). Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.
- Gx-Plus—An application that extends the 3GPP Gx interface for wireline use cases. 3GPP Gx is registered with the IANA (<http://www.iana.org>). It communicates with a PCRF.

If data for a particular AVP included in a message is not available to the router, Gx-Plus simply omits the AVP from the message it sends to the PCRF. If the PCRF determines it has insufficient information to make a determination, it may deny the request. The Diameter answer messages include the Result-Code AVP (AVP 268); the values of this AVP convey success, failure, or errors to the requestor.

- NASREQ—A Diameter-based authentication, authorization, and accounting protocol defined in RFC 7155. Junos OS supports authentication and authorization only.

Juniper Networks has also registered the Juniper-Session-Recovery application (16777296) and two new command codes (8388628 for Juniper-Session-Events and 8388629 for Juniper-Session-Discovery) with the IANA (<http://www.iana.org>).

[Table 73 on page 1732](#) describes Diameter messages the applications use.

Table 73: Diameter Messages and Diameter Applications

Diameter Message	Code	Application	Description
AA-Request (AAR)	265	JSRC, NASREQ, PTSP	Request from the application to the SAE at new subscriber login or during SAE-application synchronization. The request can be one of three types: address-authorization, provisioning-request, or synchronization.
AA-Answer (AAA)	265	JSRC, NASREQ, PTSP	Response from the SAE to the application's AA-Request message.
Abort-Session-Request (ASR)	274	JSRC, NASREQ, PTSP	Request from the SAE to the application to log out a provisioned subscriber.
Abort-Session-Answer (ASA)	274	JSRC, NASREQ, PTSP	Response from the application to the SAE's ASR message. If the application sends the logout request to AAA, the ASA message includes a success notification (ACK). If the logout failed, the ASA message includes a failure notification (NAK).
Accounting-Request (ACR)	271	JSRC, PTSP	Request from the SAE to the application or from the application to the SAE for statistics.
Accounting-Answer (ACA)	271	JSRC, PTSP	Response to the ACR message to provide statistics for each installed policy (service).

Table 73: Diameter Messages and Diameter Applications (Continued)

Diameter Message	Code	Application	Description
Capability Exchange Request (CER)	257	Gx-Plus	Request from one peer to another when the peers establish a transport connection; initiates the capability negotiation. The CER announces the peer's identity and capabilities, such as applications and security mechanisms supported.
Capability Exchange Answer (CEA)	257	Gx-Plus	Response to the CER message to announce this peer's capabilities. If this peer has no capabilities in common with the peer that sent the CER, then it must set the Result-Code AVP to DIAMETER_NO_COMMON_APPLICATION and should drop the connection. Otherwise, the CEA details establish common capabilities between the peers and enable them to further establish communication.

Table 73: Diameter Messages and Diameter Applications (Continued)

Diameter Message	Code	Application	Description
Credit-Control-Request (CCR)	272	Gx-Plus	<p>Request from Gx-Plus to the PCRF at subscriber login, logout, or update.</p> <p>An initial request (CCR-I) is sent when a subscriber logs in and AAA is requested to activate the subscriber's session. Gx-Plus retries the CCR-I message if a CCA-I message is not received from the PCRF within 10 seconds. The CCR-I message is retried up to 3 times.</p> <p>The CCR-I message includes the Diameter AVP Subscription-Id attribute (443) with the Subscription-Id-Type Diameter AVP sub-attribute (450) set to 4 (END_USER_PRIVATE) and the Subscription-Id-Data Diameter AVP sub-attribute (444) set to reserved.</p> <p>If no CCA-I is received after the 4 CCR-I messages have been sent—the first message plus 3 retries—then Gx-Plus starts sending CCR-N messages. CCR-N messages are retried forever until a success or failure response is received from the PCRF. CCR-N messages include the Juniper-Provisioning-Source AVP (AVP code 2101) set to local to notify the PCRF that the router has the authority to make a local decision regarding subscriber service activation.</p> <p>An update request (CCR-U) message is sent when a usage threshold is reached. The CCR-U reports the actual usage for all</p>

Table 73: Diameter Messages and Diameter Applications *(Continued)*

Diameter Message	Code	Application	Description
			<p>statistics. The PCRF may return a CCA-U message that includes new monitoring thresholds, service activations, service deactivations.</p> <p>If the PCRF times out on the CCR-U report, the router sets the threshold default to 10 minutes. When the change in threshold values is less than the minimum, the values are adjusted to the minimums. For example, the minimum increase for duration is 10 minutes.</p> <p>A CCR-U is also sent to report the status of service activation or deactivation. When a monitored service is deactivated separate from a subscriber logout, the CCR-U indicates that the service is no longer active and includes the service's usage data.</p> <p>A termination request (CCR-T) is sent at subscriber logout to inform the PCRF that a provisioned subscriber session is being terminated. CCR-T messages are retried forever until a success response is received from the PCRF.</p> <p>When a monitored service is deactivated as part of the subscriber logout, the CCR-T message includes monitored usage data for the service, such as bytes used.</p>

Table 73: Diameter Messages and Diameter Applications *(Continued)*

Diameter Message	Code	Application	Description
Credit-Control-Answer (CCA)	272	Gx-Plus	<p>Reply from the PCRF to a CCR message.</p> <p>In response to a CCR-I, the PCRF returns a CCA-I message that indicates success (DIAMETER_SUCCESS) or failure (DIAMETER_AUTHORIZATION_REJECTED) depending on whether the subscriber has sufficient credit for the requested services. All other responses are ignored and the CCR-I is retried.</p> <p>In response to a CCR-T, the PCRF returns a CCA-T message that indicates a successful termination with a value of 2001 (DIAMETER_SUCCESS) in the Result-Code AVP. All other responses are ignored and the CCR-T is retried.</p> <p>A CCA-N is a response to a CCR-N.</p>
Juniper-Session-Discovery-Request (JSDR)	8388629	Gx-Plus	Discovery request from the PCRF to Gx-Plus to discover subscriber sessions on the router.

Table 73: Diameter Messages and Diameter Applications (Continued)

Diameter Message	Code	Application	Description
Juniper-Session-Discovery-Answer (JSDA)	8388629	Gx-Plus	<p>Reply from router to a JSDR message; describes session information. The Result-Code AVP includes one of the following values, or an error value:</p> <ul style="list-style-type: none"> • 2001—DIAMETER_SUCCESS; the end of the database was reached, meaning all information has been sent. • 2002—DIAMETER_LIMITED_SUCCESS; some of the session information was sent, but more remains to be sent.
Juniper-Session-Event-Request (JSER)	8388628	Gx-Plus	<p>Request from router to PCRF regarding events that take place on the router. Notifies the PCRF of certain events on the router by including the Juniper-Event-Type AVP (AVP code 2103). Events reported include cold or warm boots, explicit discovery requests, substantial configuration changes, non-response or error response from PCRF, and exhaustion of fault-tolerant resources.</p>
Juniper-Session-Event-Answer (JSEA)	8388628	Gx-Plus	<p>Reply from PCRF to a JSER message.</p>
Push-Profile-Request (PPR)	288	JSRC, PTSP	<p>Request from the SAE to the router to activate or deactivate services for a subscriber.</p>

Table 73: Diameter Messages and Diameter Applications (Continued)

Diameter Message	Code	Application	Description
Push-Profile-Answer (PPA)	288	JSRC, PTSP	Response from the router to the SAE's PPR message. Includes success or failure notification for each of the service activation or deactivation commands in the request.
Re-Auth-Request (RAR)	258	Gx-Plus	<p>Audit request from the PCRF to router to determine whether a specific subscriber is still present.</p> <p>The router updates the monitoring key and threshold values when they are received in the RAR.</p>
Re-Auth-Answer (RAA)	258	Gx-Plus	<p>Reply from router to a RAR message; indicates whether the subscriber is active. The Result-Code AVP includes one of the following values:</p> <ul style="list-style-type: none"> • 2001—DIAMETER_SUCCESS; subscriber entry was found. • 5002—DIAMETER_UNKNOWN_SESSION_ID; subscriber entry was not found. • 3002—DIAMETER_UNABLE_TO_DELIVER; Gx-Plus is not configured.
Session-Resource-Query (SRQ)	277	JSRC, PTSP	Request from the router to the SAE or from the SAE to the router to initiate synchronization between router and the SAE.

Table 73: Diameter Messages and Diameter Applications (Continued)

Diameter Message	Code	Application	Description
Session-Resource-Reply (SRR)	277	JSRC, PTSP	Response to the SRQ message to begin synchronization.
Session-Termination-Request (STR)	275	JSRC, NASREQ, PTSP	Notification from the router to the SAE that a provisioned subscriber has logged out.
Session-Termination-Answer (STA)	275	JSRC, NASREQ, PTSP	Response from the SAE to the router's STR message. Includes success or failure notification.

Diameter AVPs and Diameter Applications

Diameter conveys information by including various attribute-value pairs (AVPs) in Diameter messages, in the same way that RADIUS conveys information in both standard IETF RADIUS attributes and vendor-specific attributes (VSAs). [Table 74 on page 1739](#) lists the standard Diameter AVPs used in interactions with the supported Diameter applications. Diameter reserves AVP attribute numbers 0 through 255 for RADIUS attributes that are implemented in Diameter; the Diameter attribute numbers are the same as for the corresponding standard RADIUS attributes. Attributes numbered higher than 255 have no corresponding standard RADIUS attribute. Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Table 74: Standard Diameter AVPs

Attribute Number	Diameter AVP	Application	Description	Type
1	User-Name	Gx-Plus, JSRC, NASREQ	Specifies the username. For a subscriber managed by AAA, the value is the subscriber's login name. For a static interface, the value is the interface name, which is used as the subscriber's login name.	UTF8String

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
2	User-Password	NASREQ	Specifies the password of the user to be authenticated or the user's input in a multi-round authentication exchange.	OctetString
4	NAS-IP-Address	NASREQ	Specifies the IP address of the NAS that is authenticating the user.	IPAddress
6	Service-Type	NASREQ	Specifies the type of service the user has requested or the type of service to be provided. One such AVP may be present in an authentication or authorization request or response. A NAS is not required to implement all of these service types.	Enumerated
8	Framed-IP-Address	Gx-Plus, JSRC, NASREQ, PTSP	Identifies the IPv4 address configured for the subscriber. This is the same value as for RADIUS Framed-IP-Address attribute [8].	OctetString
9	Framed-IP-Netmask	NASREQ	Identifies the four octets of the IPv4 netmask.	OctetString
11	Filter-ID	NASREQ	Specifies the name of the filter list for a user. It is intended to be human readable. Zero or more Filter-Id AVPs may be sent in an authorization answer message.	UTF8String
12	Framed-MTU	NASREQ	Specifies the maximum transmission unit (MTU) to be configured for the user, when it is not negotiated by some other means (such as PPP).	Unsigned32

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
22	Framed-Route	NASREQ	Specifies the 7-bit US-ASCII routing information.	UTF8String
25	Class	NASREQ	Returns state information from a Diameter server to the access device.	OctetString
27	Session-Timeout	NASREQ	Specifies the maximum number of seconds of service provided to the user before termination of the session.	Unsigned32
28	Idle-Timeout	NASREQ	Specifies the maximum number of consecutive seconds of idle connection allowable to the user before termination of the session or before a prompt is issued.	Unsigned32
32	NAS-Identifier	NASREQ	Specifies the identity of the NAS that provides service to the user.	DiamIdent
44	Acct-Session-ID	NASREQ	Specifies the contents of the RADIUS Acct-Session-Id attribute.	OctetString
50	Acct-Multi-Session-ID	NASREQ	Links multiple related accounting sessions, where each session has a unique Session-Id but the same Acct-Multi-Session-Id AVP.	UTF8String
55	Event-Timestamp	Gx-Plus, JSRC, PTSP	Specifies the time of the event that triggered the message in which this AVP is included. Time is indicated in seconds since January 1, 1900, 00:00 UTC.	Time

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
60	CHAP-Challenge	NASREQ	Specifies the PPP Challenge-Handshake Authentication Protocol (CHAP) challenge sent by the NAS to the CHAP peer.	OctetString
61	NAS-Port-Type	NASREQ	Specifies the type of the port on which the NAS is authenticating the user.	Enumerated
62	Port-Limit	NASREQ	Specifies the maximum number of ports the NAS provides to the user.	Unsigned32
78	Configuration-Token	NASREQ	Indicates the type of user profile used.	OctetString
85	Acct-Interim-Interval	JSRC, PTSP	<p>Specifies the number of seconds between each interim accounting update for this session.</p> <p>The router uses the following guidelines for interim accounting:</p> <ul style="list-style-type: none"> • Attribute value is within the acceptable range (600 through 86,400 seconds)—Accounting is updated at the specified interval. • Attribute value is less than the minimum acceptable value—Accounting is updated at the minimum interval (600 seconds). • Attribute value is greater than the maximum acceptable value—Accounting is updated at the maximum interval (86,400 seconds). 	Unsigned32

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
87	NAS-Port-Id	Gx-Plus, JSRC, NASREQ, PTSP	Identifies the port of the NAS that authenticates the user. This is the same value as for RADIUS NAS-Port-Id attribute [87].	UTF8String
88	Framed-Pool	NASREQ	Specifies the name of an assigned address pool to use to assign an address for the user. If a NAS does not support multiple address pools, the NAS disregards this AVP. Address pools are usually used for IP addresses but can be used for other protocols if the NAS supports pools for those protocols.	OctetString
97	Framed-IPv6-Prefix	NASREQ	Specifies the IPv6 prefix configured for the user.	OctetString
99	Framed-IPv6-Route	NASREQ	Specifies the US-ASCII routing information configured for the user on the NAS.	UTF8String
100	Framed-IPv6-Pool	NASREQ	Specifies the name of an assigned pool to use to assign an IPv6 prefix for the user. If the access device does not support multiple prefix pools, it must disregard this AVP.	OctetString
258	Auth-Application-ID	NASREQ	Specifies support of the Authentication and Authorization portion of an application.	Unsigned32
263	Session-ID	Gx-Plus, JSRC, NASREQ, PTSP	Specifies the subscriber session identifier. The router assigns the value to uniquely identify a subscriber session.	UTF8String

Table 74: Standard Diameter AVPs *(Continued)*

Attribute Number	Diameter AVP	Application	Description	Type
264	Origin-Host	NASREQ	Specifies the host that originates a Diameter message.	DiamIdent

Table 74: Standard Diameter AVPs (*Continued*)

Attribute Number	Diameter AVP	Application	Description	Type
268	Result-Code	Gx-Plus, JSRC, NASREQ, PTSP	<p>Indicates whether a request completed successfully. Provides an error code if the request failed.</p> <p>The following classes are recognized by Diameter:</p> <ul style="list-style-type: none"> • 1xxx—Informational • 2xxx—Success • 3xxx—Protocol errors • 4xxx—Transient errors • 5xxx—Permanent failures <p>Unrecognized classes, which begin with numerals 6–9 or 0, are handled as permanent failures.</p> <p>JSRC and PTSP support the following values; all non-success values are treated as permanent failures:</p> <ul style="list-style-type: none"> • 1001—DIAMETER MULTI ROUND AUTH • 2001—DIAMETER SUCCESS • 5002—DIAMETER UNKNOWN SESSION ID • 5012—DIAMETER UNABLE TO COMPLY <p>JSRC also supports the following value, which is treated as a permanent failure:</p>	Unsigned32

Table 74: Standard Diameter AVPs (*Continued*)

Attribute Number	Diameter AVP	Application	Description	Type
			<ul style="list-style-type: none"> • 3004—DIAMETER TOO BUSY; this is a transient condition, typically when the router already has a request in process for a specified subscriber. <p>Gx-Plus supports the following values for errors in a PCRF response; when these values are received or the response is malformed or unrecognizable, the request is retried.</p> <ul style="list-style-type: none"> • 3001—DIAMETER COMMAND NOT SUPPORTED; the application is not running or the command is not recognized. • 3004—DIAMETER TOO BUSY; the received message is above either the quota of downstream transactions or the outstanding message memory limit for messages from the network. • 5012—DIAMETER UNABLE TO COMPLY; the received message is greater than the local limit. 	

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
269	Product-Name	Gx-Plus	<p>Specifies the value for the Product-Name field in Capability Exchange Request (CER) and Capability Exchange Answer (CEA) messages. The value is always JUNOS unless a different name is configured with the product-name option at the [edit diameter] hierarchy level.</p> <p>If you change the product name, the router disconnects all existing connections to Diameter peers and reconnects using the new name.</p>	UTF8String
277	Auth-Session-State	JSRC, NASREQ, PTSP	<p>Indicates whether AAA session state is maintained.</p> <ul style="list-style-type: none"> 0—STATE MAINTAINED 1—NO STATE MAINTAINED 	Enumerated
279	Failed-AVP	NASREQ	<p>Specifies debugging information in cases where a request is rejected or not fully processed due to erroneous information in a specific AVP. The value of the Result-Code AVP provides information on the reason for the Failed-AVP AVP.</p>	Grouped
281	Error-Message	NASREQ	<p>Specifies a human-readable error message that may accompany a Result-Code AVP. The Error-Message AVP is not intended to be useful in real-time; do not expect network entities to parse the message.</p>	UTF8String

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
283	Destination-Realm	NASREQ	Specifies the Diameter realm to which the Diameter message is routed.	DiamIdent
293	Destination-Host	NASREQ	Specifies the host to which a Diameter message is routed.	DiamIdent
295	Termination-Cause	JSRC, NASREQ, PTSP	<p>Indicates the reason why a session was terminated on the access device.</p> <ul style="list-style-type: none"> • 1—DIAMETER LOGOUT • 2—DIAMETER SERVICE NOT PROVIDED • 3—DIAMETER BAD ANSWER • 4—DIAMETER ADMINISTRATIVE • 5—DIAMETER LINK BROKEN • 6—DIAMETER AUTH EXPIRED • 7— DIAMETER USER MOVED • 8—DIAMETER SESSION TIMEOUT 	Enumerated
296	Origin-Realm	NASREQ	Identifies the Diameter realm of the originator of a Diameter message.	DiamIdent
402	CHAP-Auth	NASREQ	Specifies the information necessary to authenticate a user using CHAP.	Grouped

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
415	CC-Request-Number	Gx-Plus	Identifies a request within a session. The combination of Session-Id and CC-Request-Type is globally unique. The number is incremented for each request during the course of a session. The number is reset when a router high availability event takes place.	Unsigned32
416	CC-Request-Type	Gx-Plus	Specifies the type of credit control request: <ul style="list-style-type: none"> • INITIAL REQUEST (1) • UPDATE REQUEST (2) • TERMINATION_REQUEST (3) • EVENT REQUEST (4) 	Enumerated
431	Granted-Service-Unit	Gx-Plus	Contains the amount that can be provided of one or more of the following requested units specified by the client: CC-Input-Octets, CC-Output-Octets, CC-Time, or CC-Total-Octets. Included in CCA-I messages, and may be included in CCA-U messages.	Grouped

Table 74: Standard Diameter AVPs (*Continued*)

Attribute Number	Diameter AVP	Application	Description	Type
443	Subscription-Id	Gx-Plus	<p>Contains the following sub-attributes that do not appear alone:</p> <ul style="list-style-type: none"> Subscription-Id-Type—(450) This subattribute has one of the following integer values: <ul style="list-style-type: none"> 0 = END_USER_E164 1 = END_USER_IMSI 2 = END_USER_SIP_URI 3 = END_USER_NAI 4 = END_USER_PRIVATE Subscription-Id-Data—(444) This sub-attribute has a value of reserved. 	Grouped
446	Used-Service-Unit	Gx-Plus	<p>Contains the amount of the requested units that have been actually used; measured from 4 when the service is activated. The units are one or more of the following requested units specified by the client: CC-Input-Octets, CC-Output-Octets, CC-Time, or CC-Total-Octets. Included in CCR-U messages.</p>	Grouped

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
480	Accounting-Record-Type	JSRC, PTSP	<p>Specifies the type of account record for service accounting:</p> <ul style="list-style-type: none"> • INTERIM_RECORD—Accounting record sent between the start and stop records, at intervals specified by the Acct-Interim-Interval AVP (AVP code 85). It contains cumulative accounting data for the existing accounting session. • START_RECORD—Accounting record sent when the service is activated to initiate the accounting session. It contains accounting data relevant to the initiation of that session. • STOP_RECORD—Accounting record sent when the service is deactivated to terminate the accounting session. It contains cumulative data relevant to that session. 	Enumerated
1001	Charging-Rule-Install	Gx-Plus, NASREQ	Requests the installation of the rule (activation of the service) designated by the included Charging-Rule-Name AVP (1005). This AVP has a vendor ID of 10415 (3GPP).	Grouped
1002	Charging-Rule-Remove	Gx-Plus	Requests the removal of the rule (deactivation of the service) designated by the included Charging-Rule-Name AVP (1005). This AVP has a vendor ID of 10415 (3GPP).	Grouped

Table 74: Standard Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
1005	Charging-Rule-Name	Gx-Plus, NASREQ	Specifies the name of a specific rule that has been installed, modified, or removed.	OctetString
1066	Monitoring-Key	Gx-Plus	Specifies which of the monitoring structures to use. Included in Charging-Rule-Install AVP (1001). The MX router does not support aggregation of statistics across services, so the value of this AVP must be different for each service. This AVP has a vendor ID of 10415 (3GPP).	OctetString
1067	Usage-Monitoring-Information	Gx-Plus	Sets monitoring thresholds. When service statistics match at least one of the granted service values, the router sends a CCR-U report with the current statistics to the PCRF. Includes the Monitoring-Key AVP (1066) and the Granted-Service-Unit AVP (431). This AVP has a vendor ID of 10415 (3GPP).	Grouped

Juniper Networks AVPs are used in addition to the standard Diameter AVPs. These AVPs have a vendor ID (enterprise number) of 2636 or 4874, and are similar in concept to RADIUS vendor-specific attributes (VSAs). [Table 75 on page 1753](#) lists the Juniper Networks AVPs that the supported Diameter applications use.

Table 75: Juniper Networks Diameter AVPs

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
213	Interface-Set-Targeting-Weight	4874	NASREQ	Specify a weight for an interface set to associate it and its member links with an aggregated Ethernet member link for targeted distribution.	Unsigned32
214	Interface-Targeting-Weight	4874	NASREQ	Specify a weight for an interface to associate it with an interface set and thus with the set's aggregated Ethernet member link for targeted distribution. When an interface set does not have a weight, then the interface weight value for the first authorized subscriber interface is used for the set.	Unsigned32
2004	Juniper-Service-Bundle	2636	JSRC	Specifies the name of the service bundle.	OctetString
2010	Juniper-DHCP-Options	2636	JSRC	Specifies the client's DHCP options.	OctetString
2011	Juniper-DHCP-GI-Address	2636	JSRC	Specifies the DHCP relay agent's IP address.	OctetString
2020	Juniper-Policy-Install	2636	JSRC, PTSP	Specifies policies to be activated for the subscriber. Includes Juniper-Policy-Name and Juniper-Policy-Definition	Grouped
2021	Juniper-Policy-Name	2636	JSRC, PTSP	Defines the name of a policy decision.	OctetString
2022	Juniper-Policy-Definition	2636	JSRC, PTSP	Defines a policy decision. Includes Juniper-Policy-Name, Juniper-Template-Name, and Juniper-Substitution.	Grouped

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2023	Juniper-Template-Name	2636	JSRC, PTSP	Specifies the profile name defined by the router. PTSP supports only the __svc_rule__ policy template.	UTF8String
2024	Juniper-Substitution	2636	JSRC, PTSP	Defines the substitution attributes. Includes Juniper-Substitution-Name and Juniper-Substitution-Value.	OctetString
2025	Juniper-Substitution-Name	2636	JSRC, PTSP	Defines the name of the variable to be replaced.	OctetString
2026	Juniper-Substitution-Value	2636	JSRC, PTSP	Defines the value of the variable to be replaced.	OctetString
2027	Juniper-Policy-Remove	2636	JSRC, PTSP	Specifies policies to be deactivated for the subscriber. Includes Juniper-Policy-Name.	Grouped
2035	Juniper-Policy-Failed	2636	JSRC, PTSP	Specifies the name of the policy activation or deactivation that failed.	OctetString
2038	Juniper-Policy-Success	2636	JSRC, PTSP	Specifies the name of the policy activation or deactivation that succeeded.	OctetString
2046	Juniper-Logical-System	2636	JSRC, PTSP	Specifies the logical system.	UTF8String
2047	Juniper-Routing-Instance	2636	JSRC, PTSP	Specifies the routing instance.	UTF8String

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2048	Juniper-Jsrc-Partition	2636	JSRC, PTSP	Specifies the logical system and routing instance for the subscriber or request. Includes Juniper-Logical-System and Juniper-Routing-Instance	Grouped
2050	Juniper-Request-Type	2636	JSRC, PTSP	Describes the type of request: <ul style="list-style-type: none"> • 1—ADDRESS_AUTHORIZATION • 2—PROVISIONING_REQUEST • 3—SYNCHRONIZATION • 4—NETWORK_FAMILY_ACTIVATE JSRC only. • 5— NETWORK_FAMILY_DEACTIVATE JSRC only. 	Enumerated
2051	Juniper-Synchronization-Type	2636	JSRC, PTSP	Describes the type of synchronization: <ul style="list-style-type: none"> • 1—FULL-SYNC • 2—FAST-SYNC • 3—NO-STATE-TO-SYNC 	Enumerated
2052	Juniper-Synchronization	2636	JSRC, PTSP	Describes the state of synchronization: <ul style="list-style-type: none"> • 1—NO-SYNC; this is the default state • 2—SYNC-IN-PROGRESS • 3—SYNC-COMPLETE 	Enumerated

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2053	Juniper-Acct-Record	2636	JSRC, PTSP	Specifies the statistics data for each policy installed for this subscriber. Includes Juniper-Policy-Name.	Grouped
2054	Juniper-Acct-Collect	2636	JSRC, PTSP	<p>Specifies whether to collect accounting data for the installed policy (service) when included in the Juniper-Policy-Install AVP:</p> <ul style="list-style-type: none"> • 1—COLLECT_ACCT • 2—NOT_COLLECT_ACCT 	Enumerated
2058	Juniper-State-ID	2636	JSRC, PTSP	<p>Specifies the value assigned to each synchronization cycle for the purpose of identifying which messages to discard. All solicited requests containing the same Juniper-State-ID belong to the same Session-Resource-Query (SRQ) synchronization cycle. Messages from a previous synchronization cycle are discarded. When a new cycle begins, the value of the Juniper-State-ID AVP is increased by 1.</p> <p>NOTE: For solicited synchronization requests, the SRQ message contains the incremented Juniper-State-ID value. For unsolicited synchronization requests, the Session-Resource-Reply (SRR) message contains the incremented Juniper-State-ID value.</p>	Unsigned32
2100	Juniper-Virtual-Router	2636	Gx-Plus, JSRC	Specifies the name of the virtual router associated with the session.	UTF8String

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2101	Juniper-Provisioning-Source	2636	Gx-Plus	<p>Specifies the provisioning source for the session in CCR-N and JSDA messages:</p> <ul style="list-style-type: none"> • 1—Local • 2—Remote 	Enumerated
2102	Juniper-Provisioning-Descriptor	2636	Gx-Plus	<p>Defines the group used in JSDA messages that includes the session ID, and optionally Juniper-Provisioning-Source and subscriber data.</p>	Grouped
2103	Juniper-Event-Type	2636	Gx-Plus	<p>Communicates the event type in JSER messages:</p> <ul style="list-style-type: none"> • 1—Cold boot; all sessions are lost • 2—Warm boot; sessions are preserved • 3—Discovery requested by the operator • 4—<i>Are you there?</i> (AYT); application level ping sent when the notification is due to no response or an erroneous response from the PCRF, or due to a configuration change. • 5—AWD; application-level watchdog sent by the router when there has been no other activity for 15 seconds. The watchdog is sent every 5 seconds unless preempted by higher-priority synchronization event. 	Enumerated

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2104	Juniper-Discovery-Descriptor	2636	Gx-Plus	Defines the group used in JS DR and JS DA messages that includes parameters of a discovery request: discovery type, request string, verbosity, max results.	Grouped
2105	Juniper-Discovery-Type	2636	Gx-Plus	Specifies the discovery subcommand for JS DR and JS DA messages: <ul style="list-style-type: none"> • 1—Exact: look up the data for the specified session. • 2—Bulk: Provide get-bulk kinds of information after the specified string. • 3—Done: Stop retries for all sessions up to the specified session. 	Enumerated
2106	Juniper-Verbosity-Level	2636	Gx-Plus	Specifies the verbosity level for JS DR and JS DA messages: <ul style="list-style-type: none"> • 1—Summary; include only the Session-Id AVP. • 2—Brief; include the Session-Id, Juniper-Virtual-Router, and Framed-IP-Address AVPs. • 3—Detail; include the Session-Id, Juniper-Provisioning-Source, Juniper-Virtual-Router, Framed-IP-Address, and Event-Timestamp AVPs. • 4—Extensive; include all available session information. 	Enumerated
2107	Juniper-String-A	2636	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2108	Juniper-String-B	2636	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String
2109	Juniper-String-C	2636	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String
2110	Juniper-Unsigned32-A	2636	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32
2111	Juniper-Unsigned32-B	2636	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32
2112	Juniper-Unsigned32-C	2636	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32
2200	Juniper-IPv6-Ndra-Prefix	2636	JSRC	<p>If available in the subscriber's session database IPv6Prefix entry, this AVP is included in AAR provisioning request messages sent to the SAE.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	IPv6Prefix
2201	Juniper-Framed-IPv6-Netmask	2636	JSRC	<p>If available in the subscriber's session database IPv6Address entry, this AVP is included in AAR provisioning request messages sent to the SAE.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	IPv6Address

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2202	Juniper-Agent-Circuit-Id	2636	JSRC	<p>Identifies the subscriber by access node and subscriber line. If available in the subscriber's session database entry, this AVP is included in AAR provisioning request messages sent to the SAE.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	OctetString
2203	Juniper-Agent-Remote-Id	2636	JSRC	<p>Identifies the subscriber on the access node. If available in the subscriber's session database entry, this AVP is included in AAR provisioning request messages sent to the SAE.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	OctetString
2204	Juniper-Acct-IPv6-Input-Octets	2636	JSRC	<p>Number of IPv6 octets received on the interface. This AVP is included in ACR accounting request messages sent to the SAE, even when the value is zero.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	Unsigned64
2205	Juniper-Acct-IPv6-Output-Octets	2636	JSRC	<p>Number of IPv6 octets sent on the interface. This AVP is included in ACR accounting request messages sent to the SAE, even when the value is zero.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	Unsigned64

Table 75: Juniper Networks Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2206	Juniper-Acct-IPv6-Input-Pkts	2636	JSRC	<p>Number of IPv6 packets received on the interface. This AVP is included in ACR accounting request messages sent to the SAE, even when the value is zero.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	Unsigned64
2207	Juniper-Acct-IPv6-Output-Pkts	2636	JSRC	<p>Number of IPv6 packets sent on the interface. This AVP is included in ACR accounting request messages sent to the SAE, even when the value is zero.</p> <p>This AVP is used only when you enable JSRC dual-stack support.</p>	Unsigned64

Telelec AVPs are used only for Gx-Plus. These AVPs have an enterprise number of 21274. [Table 76 on page 1761](#) lists the Telelec AVPs. These four variables are used to provide substitution values for user-defined CoS service variables.

Table 76: Telelec Diameter AVPs

Attribute Number	Diameter AVP	Application	Description	Type
5555	Telelec-Charging-Rule-Argument-Name	Gx-Plus	Defines the name of the service variable to be replaced.	OctetString
5556	Telelec-Charging-Rule-Argument-Value	Gx-Plus	Defines the value of the service variable to be replaced.	OctetString

Table 76: Tekelec Diameter AVPs (Continued)

Attribute Number	Diameter AVP	Application	Description	Type
5557	Tekelec-Charging-Rule-Argument	Gx-Plus	Defines the substitution attributes used to replace service variables. Includes Tekelec-Charging-Rule-Argument-Name AVP (5555) and Tekelec-Charging-Rule-Argument-Value AVP (5556).	Grouped
5558	Tekelec-Charging-Rule-With-Arguments	Gx-Plus	Requests the installation of the rule (activation of the service) designated by the included Charging-Rule-Name AVP (1005). Requested service variable substitutions are provided by the optionally included Tekelec-Charging-Rule-Argument AVP (5557).	Grouped

Configuring Diameter

You configure Diameter by specifying the endpoint origin, the remote peers, the transport layer connection, and network elements that associate routes with peers. Only the master Diameter instance is currently supported. You can configure alternative values for this Diameter instance only in the context of the default routing instance.

To configure Diameter base protocol:

1. Configure the origin realm and origin host of the Diameter instance.
See ["Configuring the Origin Attributes of the Diameter Instance" on page 1763](#)
2. Configure the Diameter peers.
See ["Configuring Diameter Peers" on page 1763](#)
3. (Optional) Configure the Diameter transport layer elements.
See ["Configuring the Diameter Transport" on page 1765](#)
4. (Optional) Configure the Diameter network elements.
See ["Configuring Diameter Network Elements" on page 1766](#)
5. (Optional) Configure trace options for troubleshooting the configuration.
See [Tracing Diameter Base Protocol Processes for Subscriber Access](#).

Configuring the Origin Attributes of the Diameter Instance

You can configure the identifying characteristics of the endpoint node that originates Diameter messages for the Diameter instance. The hostname is supplied as the value for the Origin-Host AVP by the Diameter instance. The realm is supplied as the value for the Origin-Realm AVP by the Diameter instance.

To configure the origin attributes for a Diameter instance:

1. Specify the name of the host that originates the Diameter message.

```
[edit diameter origin]
user@host# set host host14
```

2. Specify the realm of the host that originates the Diameter message.

```
[edit diameter origin]
user@host# set realm example.com
```

Configuring Diameter Peers

You can configure the peers to which Diameter sends messages. Diameter uses the default logical system and routing instance. Port 3868 is used for active connections to peers by default.

To configure a remote peer for a Diameter instance:

1. Specify the name of the Diameter peer.

```
[edit diameter]
user@host# edit peer peer-name
```

2. Specify the IP address of the Diameter peer. Starting in Junos OS Release 17.3R1, both IPv4 and IPv6 addresses are supported.



NOTE: You must configure the same address family type for the peer and the corresponding local Diameter transport connection.

```
[edit diameter peer peer-name]
user@host# set address ip-address
```

3. (Optional) Specify a routing instance, a logical system, or a logical system and routing instance for the Diameter peer.

```
[edit diameter peer peer-name]
user@host# set routing-instance routing-instance-name
```

```
[edit diameter peer peer-name]
user@host# set logical-system logical-system-name
```

```
[edit diameter peer peer-name]
user@host# set logical-system logical-system-name routing-instance routing-instance-name
```

4. (Optional) Specify the port that Diameter uses for active connections to the peer.

```
[edit diameter peer peer-name]
user@host# set connect-actively port port-number
```

5. (Optional) Specify the transport that Diameter uses for active connections to the peer.

```
[edit diameter peer peer-name]
user@host# set connect-actively transport transport-name
```

6. (Optional) Specify the name of the peer host and the name of the peer realm.



NOTE: You must specify both the host and realm for the peer origin.

```
[edit diameter peer peer-name]
user@host# set peer-origin host hostname realm realm-name
```

7. (Optional) Include the Origin-State attribute-value pair (AVP) for the Diameter peer in Diameter base protocol-level messages to enable monitoring of changes in the AVP value.

```
[edit diameter peer peer-name]
user@host# set send-origin-state-id
```

For example, the following configuration for peer p3 specifies an IPv4 address, the routing instance ri8, destination port 49152, transport t6, an origin of host 1 in example.com, and includes the Origin-State AVP in messages.

```
[edit diameter]
user@host# edit peer p3
[edit diameter peer p3]
user@host# set address 192.168.23.10
user@host# set routing-instance ri8
user@host# set connect-actively port 49152
user@host# set connect-actively transport t6
user@host# set peer-origin host host1 realm example.com
user@host# set send-origin-state-id
```

Configuring the Diameter Transport

You can configure one or more transports for a Diameter instance to set the IPv4 or IPv6 address for the local connection, and optionally configure a logical system or routing instance context. Diameter uses the default logical system and routing instance. The logical system and routing instance for the transport connection must match that for the peer, or a configuration error is reported. Multiple peers can share the same transport.

To configure a transport for a Diameter instance:

1. Configure the transport name.

```
[edit diameter]
user@host# edit transport transport-name
```

2. Configure the local IP address for the Diameter local transport connection. Starting in Junos OS Release 17.3R1, both IPv4 and IPv6 addresses are supported.



NOTE: The address family must match that for the remote Diameter peer.

```
[edit diameter transport t1]
user@host# set address ip-address
```

3. (Optional) Configure a logical system and optionally a routing instance for the transport.

```
[edit diameter transport t1]
user@host# set logical-system logical-system-name
```

4. (Optional) Configure a routing instance for the transport.

```
[edit diameter transport t1]
user@host# set routing-instance routing-instance-name
```

For example, the following configuration for transport t1 specifies an IPv6 address, logical system ls5, and routing instance ri10.

```
[edit diameter]
user@host# edit transport t1
[edit diameter transport t1]
user@host# set address 2001:db8::113:200
user@host# set logical-system ls5
user@host# set routing-instance ri10
```

Configuring Diameter Network Elements

A Diameter network element (DNE) consists of associated applications (called functions in the CLI), a list of prioritized peers, and a set of forwarding rules. The forwarding rules define individual routes through a set of associated destinations, applications, and metrics. At least one DNE must be configured per chassis to start the Diameter process (jdiameterd).

Before you configure Diameter network elements, perform the following task:

- Define the Diameter peers. See ["Configuring Diameter Peers" on page 1763](#).

To configure a Diameter network element:

1. Specify the name of the network element.

```
[edit diameter]
user@host# set network-element dne25
```

2. (Optional) Associate one or more applications with the network element. All applications are associated by default.

```
[edit diameter network-element dne25]  
user@host# set function jsrc
```

3. Associate a Diameter peer with the network element and set the priority for the peer.

```
[edit diameter network-element dne25]  
user@host# set peer peer1 priority 1
```

4. Specify a route that is reachable through the network element based on the forwarding rules that you define.

```
[edit diameter network-element dne25]  
user@host# set forwarding route dne-route2
```

5. Specify a metric for the route.

```
[edit diameter network-element dne25 forwarding route dne-route2]  
user@host# set metric 15
```

6. (Optional) Associate the route with a destination host and realm.

```
[edit diameter network-element dne25 forwarding route dne-route2]  
user@host# set destination host host5 realm example.com
```

7. (Optional) Specify an application associated with the route.

```
[edit diameter network-element dne25 forwarding route dne-route2]  
user@host# set function jsrc
```

8. (Optional) Specify the realm of the network element origin and optionally also specify the name of the element host.



NOTE: Only the realm name is required.

```
[edit diameter peer p3]  
user@host# set dne-origin realm realm-name <host hostname>
```

Example: Configure S6a Application

IN THIS SECTION

- [Requirements | 1768](#)
- [Overview | 1768](#)
- [Configuration | 1769](#)
- [Verification | 1778](#)

This example shows how to configure diameter-based authentication S6a application on your SRX Series Firewall to retrieve authentication information from the subscriber server.

Requirements

This example uses the following hardware:

- Any SRX Series Firewall

Before you begin, read "[Diameter Base Protocol Overview](#)" on page 1728.

Overview

In this example, You create S6a partition and specify the endpoint origin, the remote peers, and the network elements that associate routes with peers to control diameter forwarding of S6a messages. You also create S6a partition to Only the master Diameter instance is currently supported. You can configure alternative values for the master Diameter instance only in the context of the default routing instance.

Configuration

IN THIS SECTION

- [Configure Access Profile and Diameter Application Parameters | 1769](#)
- [Configure Redundant Ethernet Interfaces | 1773](#)
- [Configure Security Zones and Security Policies to permit the S6a Diameter Application | 1776](#)

Configure Access Profile and Diameter Application Parameters

CLI Quick Configuration

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

```
set access-profile s6a_test authentication-order s6a
set access profile s6a_test authentication-order s6a
set access s6a partition partition_name
set access s6a partition partition_name destination-realm zzz.com
set access s6a partition partition_name destination-host s6b.zzz.com
set access s6a partition partition_name diameter-instance master
set access s6a partition partition_name max-outstanding-requests 40
set access s6a partition partition_name response-timeout 20
set diameter origin realm zzz.com
set diameter origin host s6a.zzz.com
set diameter network-element ne3
set diameter network-element peer p3
set diameter network-element peer p3 priority 100
set diameter network-element ne3 forwarding route r0
set diameter network-element ne3 forwarding route r0 metric 100
set diameter peer p3 address 192.0.0.244
set diameter peer p3 connect-actively port 63101
```


Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure access profile and diameter application parameters:

1. Specify the access profile to use for authentication order.

```
[edit access-profile]
user@host# set s6a_test
```

2. Specify the order in which authentication methods are used.

```
[edit access profile]
user@host# set s6a_test authentication-order s6a
```

3. Create the partition or specify the name of an existing partition.

```
[edit access]
user@host# set s6a partition partition_name
```

4. Configure the destination realm for the s6a partition.

```
[edit access]
user@host# set s6a partition partition_name destination-realm zzz.com
```

5. Configure the destination host for the s6a partition.

```
[edit access]
user@host# set s6a partition partition_name destination-host s6b.zzz.com
```

6. Specify the Diameter instance for the s6a partition.

```
[edit access]
user@host# set s6a partition partition_name diameter-instance master
```



NOTE: Currently, only the default Diameter instance, `master`, is supported.

7. Set a limit on the number of outstanding requests.

```
[edit access]
user@host# set s6a partition partition_name max-outstanding-requests 40
```

8. Configure the amount of time in seconds before the s6a stops attempting to send a subscriber logout message.

```
[edit access]
user@host# set s6a partition partition_name response-timeout 20
```

9. Include the name of the realm that originates the Diameter message.

```
[edit diameter]
user@host# set origin realm zzz.com
```

10. Include the name of the host that originates the Diameter message.

```
[edit diameter]
user@host# set origin host s6a.zzz.com
```

11. Specify the name of the network element.

```
[edit diameter]
user@host# set network-element ne3
```

12. Associate a Diameter peer with the network element.

```
[edit diameter]
user@host# set network-element peer p3
```

13. Set the priority for the peer.

```
[edit diameter]
user@host# set network-element peer p3 priority 100
```

14. Specify a route that is reachable through the network element based on the forwarding rules that you define.

```
[edit diameter]
user@host# set network-element ne3 forwarding route r0
```

15. Specify a metric for the route.

```
[edit diameter]
user@host# set network-element ne3 forwarding route r0 metric 100
```

16. Specify the IP address of the Diameter peer.

```
[edit diameter]
user@host# set peer p3 address 192.0.0.244
```

17. Specify the port that Diameter uses for active connections to the peer.

```
[edit diameter]
user@host# set peer p3 connect-actively port 63101
```

Results

From configuration mode, confirm your configuration by entering the `show access` and `show diameter` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show access
s6a {
  partition partition_name {
    destination-realm zzz.com;
```

```

        destination-host s6b.zzz.com;
        diameter-instance master;
        max-outstanding-requests 40;
        response-timeout 20;
    }
}

```

```

[edit]
user@host# show diameter
    origin {
        realm zzz.com;
        host s6a.zzz.com;
    }
    network-element ne3 {
        forwarding {
            route r0 {
                metric 100;
            }
        }
    }

    peer p3 {
        address 192.0.0.244;
        connect-actively {
            port 63101;
        }
    }
}

```

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

Configure Redundant Ethernet Interfaces

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy

and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 gigether-options redundant-parent reth0
set interfaces ge-0/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 192.0.0.254/8
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 198.51.100.254/8
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure redundant Ethernet interfaces:

1. Configure redundant Ethernet interfaces.

```
[edit interfaces]
user@host# set ge-0/0/0 gigether-options redundant-parent reth0
user@host# set ge-0/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 192.0.0.254/8
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 198.51.100.254/8
```

Results

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

```
[edit]
user@host# show interfaces
ge-0/0/0 {
```

```

    gigether-options {
        redundant-parent reth0;
    }
}
ge-0/0/1 {
    gigether-options {
        redundant-parent reth1;
    }
}
ge-7/0/0 {
    gigether-options {
        redundant-parent reth0;
    }
}
ge-7/0/1 {
    gigether-options {
        redundant-parent reth1;
    }
}
reth0 {
    redundant-ether-options {
        redundancy-group 1;
    }
    unit 0 {
        family inet {
            address 192.0.0.254/8;
        }
    }
}
reth1 {
    redundant-ether-options {
        redundancy-group 1;
    }
    unit 0 {
        family inet {
            address 198.51.100.254/8;
        }
    }
}
}

```

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

Configure Security Zones and Security Policies to permit the S6a Diameter Application

CLI Quick Configuration

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

```
set security zones security-zone Outside host-inbound-traffic system-services all
set security zones security-zone Outside host-inbound-traffic protocols all
set security zones security-zone Outside interfaces reth1.0
set security zones security-zone Inside host-inbound-traffic system-services all
set security zones security-zone Inside host-inbound-traffic protocols all
set security zones security-zone Inside interfaces reth0.0
set security policies from-zone Inside to-zone Outside policy policy0 match source-address any
set security policies from-zone Inside to-zone Outside policy policy0 match destination-address any
set security policies from-zone Inside to-zone Outside policy policy0 match application any
set security policies from-zone Inside to-zone Outside policy policy0 then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure security policies and zones:

1. Set system services and protocols on reth1.0 interface.

```
[edit security]
user@host# set zones security-zone Outside host-inbound-traffic system-services all
user@host# set zones security-zone Outside host-inbound-traffic protocols all
user@host# set zones security-zone Outside interfaces reth1.0
```

2. Set system services and protocols on reth0.0 interface.

```
[edit security]
user@host# set zones security-zone Inside host-inbound-traffic system-services all
```

```

user@host# set zones security-zone Inside host-inbound-traffic protocols all
user@host# set zones security-zone Inside interfaces reth0.0

```

3. Configure the security policies.

```

[edit security ]
user@host# set policies from-zone Inside to-zone Outside policy policy0 match source-address
any
user@host# set policies from-zone Inside to-zone Outside policy policy0 match destination-
address any
user@host# set policies from-zone Inside to-zone Outside policy policy0 match application any
user@host# set policies from-zone Inside to-zone Outside policy policy0 then permit

```

Results

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

```

[edit]
user@host# show security policies
    from-zone Inside to-zone Outside {
        policy policy0 {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
}

```

```

[edit]
user@host# show security zones
    security-zone Outside {
        host-inbound-traffic {
            system-services {

```



```
        all;
    }
    protocols {
        all;
    }
}
interfaces {
    reth1.0;
}
}
security-zone Inside {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        reth0.0;
    }
}
```

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

Verification

IN THIS SECTION

- [Verifying the S6a Status | 1778](#)

Verifying the S6a Status

Purpose

To confirm that the configuration is working properly, perform these tasks:

Action

From operational mode, enter the `show network-access s6a state`, `show network-access s6a statistics`, and `show network-access s6a statistics extensive` commands to check the network access state and statistics of s6a application.

```
user@host> show network-access s6a state
```

S6a state:

Component	Value
active-configuration	yes
queue-state	normal
request-count	0

```
user@host> show network-access s6a statistics
```

S6a general counters:

Counter	Value
aia-grant	1

```
user@host> show network-access s6a statistics extensive
```

S6a general counters:

Counter	Value
air	0
air-retry	0
air-failures	0
aia	0
aia-grant	0
aia-deny	0
aia-timeout	0
aia-failure	0
aia-late-response	0
aia-parse-errors	0
aia-drops-no-session	0
aia-drops-bad-orealm	0
aia-drops-bad-ohost	0
aia-drops-no-result	0
aia-drops-other	0
aia-bad-result	0
aia-bad-data	0
rx-unsupported-resp-cmd	0

rx-bad-experimental-result	0
rx-bad-authentication-info	0
rx-bad-utran-vector	0
rx-bad-eutran-vector	0
rx-bad-geran-vector	0
rx-parse-errors	0

S6a diameter event counters:

Diameter event	Value
bad data message	0
good data message	0
bad flags	0
bad fixed destination	0
bad routed destination	0
tx is over limit	0
bad end-to-end id	0
no peer for tx	0
peer down while waiting for answer	0
timeout while waiting for answer	0
tx timeout	0
tx try limit	0
tx failure	0
discarded	0
received answer is over limit	0
tx failure: no memory	0
base-app-tx-timeout	0
base-app-rx-timeout	0
base-app-tx-discard	0
base-app-rx-discard	0

Meaning

The `show network-access s6a state`, `show network-access s6a statistics`, and `show network-access s6a statistics extensive` commands show the S6a application state and the statistics of the retrieved authentication information from the subscribed server.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, both IPv4 and IPv6 addresses are supported.
17.3R1	Starting in Junos OS Release 17.3R1, both IPv4 and IPv6 addresses are supported.
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

RELATED DOCUMENTATION

[Gx-Plus for Provisioning Subscribers | 1781](#)

[3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1799](#)

[NASREQ for Authentication and Authorization | 1852](#)

[JSRC for Subscriber Provisioning and Accounting | 1856](#)

Gx-Plus for Provisioning Subscribers

IN THIS SECTION

- [Gx-Plus for Provisioning Subscribers Overview | 1782](#)
- [Understanding Gx-Plus Interactions Between the Router and the PCRF | 1784](#)
- [Configuring Gx-Plus | 1793](#)
- [Configuring the Gx-Plus Partition | 1794](#)
- [Configuring Gx-Plus Global Attributes | 1795](#)
- [Provisioning Subscribers with Gx-Plus | 1796](#)

- [Disabling PCRF Control of a Subscriber Session | 1797](#)

Gx-Plus for Provisioning Subscribers Overview

IN THIS SECTION

- [Benefits of Gx-Plus | 1784](#)

Gx-Plus is a Diameter-based application that extends the capability of the Gx interface. The 3rd Generation Partnership Project (3GPP) defined Gx as the online policy interface between the Policy Control and Charging Rules Function (PCRF) and the Policy and Charging Enforcement Function (PCEF), to provide control over policy and flow-based charges for subscribers. The PCRF is a centralized policy decision point that deploys business policy rules to allocate broadband network resources and manages flow-based charges for subscribers and services. The router functions as the PCEF in this environment.

Gx-Plus provides provisioning, activation, and deactivation of services; threshold triggers for service statistics processing; service accounting; subscriber session termination; fault recovery; and event (subscriber login and logout) notifications. The terminology typically used for PCRFs varies slightly from standard Junos OS terminology. The terms listed in [Table 77 on page 1782](#) are interchangeable.

Table 77: Differences Between Gx-Plus and Junos OS Terminology

Gx-Plus	Junos OS
policy	service
rule	service
rule install or installation	service activation or instantiation
rule uninstall	service deactivation
usage monitoring	service accounting

Gx-Plus enables the router acting as a PCEF to exchange Diameter Credit-Control Application (DCCA) messages with a PCRF residing on a server to request credit authorization and service provisioning for authenticated subscribers. When an application requests AAA to activate a subscriber's session, the router sends a Credit-Control-Request (CCR) message to determine whether the subscriber has credit for the desired services and to request provisioning of those services from the PCRF policy manager.

The PCRF responds with a Credit-Control-Answer (CCA) message that indicates success or failure for credit authorization. If the subscriber has sufficient credit for the requested services, credit is authorized. If the subscriber has insufficient credit for the services, credit authorization fails.

The CCA can include services to be activated for the subscriber. If the response times out, the subscriber is logged in but only default services—if present—are activated for the subscriber. The router interprets the omission of the Result-Code AVP from the CCA as a provisioning authorization failure and does not allow the subscriber to log in.

When a subscriber client application, such as DHCP, sends a subscriber logout notice to AAA, the router in turn sends a CCR message to the PCRF to request subscriber termination. The PCRF acknowledges the logout with a CCA message.

Different Diameter message types exchanged by the router and the PCRF contain different sets of attribute-value pairs (AVPs). If data for an AVP is not available for a request to the PCRF, that AVP is omitted from the message. If the PCRF subsequently has insufficient information to decide on the request, it may deny the request.

Gx-Plus establishes sessions that correspond to IPv4 DHCP sessions on dual-stack IPv6/IPv4 or IPv4-only subscriber interfaces, depending on the access profile. By default, IPv6 information is not communicated to the PCRF. You must explicitly configure Gx-Plus to include IPv6 information. When you do so, Gx-Plus can establish sessions that correspond to DHCPv6 sessions on IPv6-enabled subscriber interfaces and on dual-stack IPv6/IPv4-enabled interfaces.

For dual-stack DHCP subscribers (DHCPv4 and DHCPv6 on the same VLAN), each DHCP session is treated as a separate Gx-Plus session. However, only a single Gx-Plus session exists for dual-stack PPP sessions.

Gx-Plus includes the following fault tolerance and recovery capabilities:

- Unlimited retries of unacknowledged provisioning requests
- Unlimited retries of logout requests
- Router event notification
- Router discovery



NOTE: More than one Diameter-based application (function), such as Gx-Plus or JSRC, can run on a router simultaneously.

Benefits of Gx-Plus

- Extends the 3GPP Gx interface to provide provisioning, activation, and deactivation of services; threshold triggers for service statistics processing; service accounting; subscriber session termination; fault recovery; and event (subscriber login and logout) notifications.

Understanding Gx-Plus Interactions Between the Router and the PCRF

IN THIS SECTION

- [Subscriber Login | 1785](#)
- [Fault Tolerance and Event Notification | 1787](#)
- [PCRF-Generated Discovery | 1789](#)
- [Subscriber Accounting | 1789](#)
- [Subscriber Usage Thresholds | 1789](#)
- [Subscriber Audit | 1793](#)
- [Subscriber Logout | 1793](#)

This topic describes the sequences of Diameter messages exchanged by means of Gx-Plus between the Policy Control and Rules Charging Function (PCRF) and the router acting as a Policy and Charging Enforcement Function (PCEF) as they interact to perform the following tasks for subscriber access:

- Subscriber login
- Fault tolerance and event notification
- Subscriber usage thresholds and monitoring
- Subscriber audit
- Subscriber logout

Subscriber Login

Gx-Plus provisioning is enabled for subscribers when you include the provisioning-order `gx-plus` statement at the `[edit access profile profile-name]` hierarchy level. When an application requests AAA to activate the subscriber's session, the router sends a CCR-I message to the PCRF to request provisioning for the subscriber session. The CCR-I message must include the Juniper-Virtual-Router, Framed-IP-Address, and NAS-Port-ID AVPs. The request is not generated when no IPv4 address has been assigned to the subscriber, when IPv6 is enabled and an IPv6 address has been assigned, or when the NAS-Port-ID is unknown. Starting in Junos OS Release 17.4R1, the CCR-I message includes the Subscription-Id AVP (AVP code 443) with the Subscription-Id-Type AVP set to 4 and Subscription-Id-Data AVP set to reserved.

The PCRF returns a CCA-I message that includes the Result-Code AVP (AVP code 268). The router considers a CCA-I that does not include the Result-Code AVP as a failed response. The CCA-I can return the Charging-Rule-Install AVP (AVP code 1001), which identifies services to be activated.

If the Result-Code value is `DIAMETER_SUCCESS` (2001), the router communicates to AAA that the requested service is activated. If the Result-Code value is `DIAMETER_AUTHORIZATION_REJECTED`, the router communicates to AAA that the service activation is not permitted. If the Result-Code AVP has any other value, or is missing, the request is retried. A total of three CCR-I messages can be sent.

If the PCRF does not indicate success or failure, then by default the router continues to send requests, but the retry requests are CCR-N messages (no-response notifications) that include the Juniper-Provisioning-Source AVP (AVP code 2101). This AVP indicates that the router has local decision-making authority to provision services in the absence of a PCRF response to the CCR-I. This AVP is not present in the CCR-I message.

A subscriber login initiates the following sequence of events:

1. A client application—such as DHCP, PPP, or static subscriber sessions—requests AAA to authenticate the subscriber.
2. Authentication begins if the subscriber access profile specifies RADIUS authentication. Login continues when the authentication is successful. Login fails when the authentication-order statement in the profile does not specify RADIUS authentication or no authentication. Login also fails when authentication fails.
3. Default services are activated for the subscriber. Any services that the authentication server includes in the authentication grant are activated. Additionally, a default service may have been configured for the client application.
4. If the subscriber access profile specifies Gx-Plus provisioning, the router initiates the Gx-Plus message exchange by sending a CCR-I message to the PCRF. The router waits for the PCRF to respond with a CCA-I message within a non-configurable timeout period.

When the PCRF responds within the timeout period and includes the Charging-Rule-Install AVP in the CCA-I message, subscriber login is delayed while the router deactivates any default services and attempts to activate the specified services.

- If all the specified services are activated, then the login completes.
- If any of the services cannot be activated, the router sends the PCRF a CCR-U message with the status of the services (a rule report). The PCRF responds to this message with a CCA-U that can contain a new set of services for activation.
- The router ignores any default services, even if the CCA-I message does not include any services. In this circumstance, no services are activated.

If the PCRF does not return a CCA-I within the timeout period, subscriber login completes.

- The router searches first for services returned from the authentication server and activates any it finds. If no such services are found, then the router activates any locally configured default services. Subscriber login completes when default service activation is successful, but fails when any default service fails to activate. Because default services are not required to be present, login also completes when no default services are found.
 - If login completes (with or without a default service), the router periodically resends the CCR-I message to the PCRF. If the PCRF subsequently returns a CCA-I, the router deactivates the default service, if any, and then activates any services included in the CCA-I. If the message does not include any services, then no service is activated, not even a default service.
 - If any of the services contained in the CCA-I cannot be activated, the router sends the PCRF a CCR-U message with the status of the services (a rule report). The PCRF responds to this message with a CCA-U that can contain a new set of services for activation.
5. The router begins to monitor session accounting statistics if the CCA-I message includes any threshold triggers for usage monitoring. The Usage-Monitoring-Information AVP (AVP code 1067) contains the threshold triggers in the Granted-Service-Unit AVP (AVP code 431). The triggers are the values granted by the PCRF for the following statistics: duration of the session, input octets count, output octets count, and total octets count.
- a. If the service statistics meet or exceed any of these trigger thresholds during the session, the router sends a CCR-U message to the PCRF with accounting information in the Usage-Monitoring-Information AVP (AVP code 1067). The AVP now contains the Used-Service-Unit AVP (AVP code 446) to report the current values for all four statistics.
 - b. In response, the PCRF may return a CCA-U message with the Usage-Monitoring-Information AVP, which can include any of the following: the Granted-Service-Unit AVP with new threshold triggers (absolute values rather than increments to the previous thresholds), the Charging-Rule-Install AVP (AVP code 1001) for service activations, or the Charging-Rule-Remove AVP (AVP code 1002) for service deactivations.



NOTE: The router does not aggregate statistics across services.

6. When the subscriber logs out, the router sends a CCR-T message (termination notice) to the PCRF, which responds with a CCA-T message.

Fault Tolerance and Event Notification

Although the probability is low, the PCRF and the router can have different values for the number of subscribers. This error can arise from the following scenarios:

- CCA-I loss: if no CCA-I is delivered to the router, then the PCRF considers a subscriber as provisioned whereas the router considers it not provisioned.
- CCR-T loss: if no CCR-T is delivered to the PCRF, then the PCRF considers a subscriber to be provisioned whereas the router considers the subscriber not provisioned (logged out).

Loss of messages can be greater during cold boots and high availability events. Unacknowledged CCR-I and CCR-T requests are retransmitted forever until a satisfactory response is received to reduce the incidence of failure, and significant events are reported to Gx-Plus. By default the number of outstanding requests is limited to:

- Maximum outstanding provisioning requests = 256k
- Maximum outstanding logout requests = 256k
- Maximum outstanding rule-reports = 100k
- Maximum outstanding threshold-reports = 100k

to avoid overloading the PCRF. This limit reduces the possibility of losing requests until the subscriber session is active. You can modify this number by including the `max-outstanding-requests` statement at the `[edit access-gx-plus global]` hierarchy level.

Gx-Plus does not rely on the connection state between devices to detect router or PCRF outages, because some events do not affect the connection state and others are not detected when there is a Diameter relay or proxy between the devices. Event notifications (JSER messages) are sent when certain events take place on the router. The Juniper-Event-Type AVP (AVP code 2103) in the message describes the event.

Event notifications are retried until Gx-Plus returns a JSEA message with a Result-Code value of `DIAMETER_SUCCESS` (2001) to acknowledge receipt of the event notification. When retrying notifications, one notification is sent for each outstanding event. No other request are sent as long as there is any outstanding event other than an application watch dog (AWD).

[Table 78 on page 1788](#) lists router events and the subsequent router and PCRF actions.

Table 78: Router Events, Router Actions, and PCRF Actions

Router Event	Router Action	PCRF Action
The router receives no response from the PCRF or an error response.	Send event notification.	Respond to event notification.
The configuration changes. Significant changes such as the origin host or realm and the Gx-Plus partition destination host or realm also increment the value of the Origin-State-Id AVP.	Send event notification.	Respond to event notification and perform discovery.
The router receives an explicit discovery request from the PCRF.	Send event notification.	Respond to event notification.
The router undergoes a cold boot and all sessions are lost. This can result from a catastrophic failure or power cycle.	Send event notification.	Respond to event notification and clear the database.
The router undergoes a warm boot.	Send event notification.	Respond to event notification and clear the database.
Recovery resources that are needed to continuously retry unacknowledged requests (CCR-N and CCR-T messages) are exhausted. The value of the Origin-State-Id AVP is incremented. This event is unlikely to occur.	Send event notification.	Respond to event notification and perform discovery.

An important aspect of Gx-Plus fault tolerance is that subscriber login and termination requests are retried (replayed) forever until a satisfactory response is received from the PCRF. In rare circumstances, this can result in a stack of pending requests being replayed over and over.

You can issue the `clear network-access gx-plus replay` command to clear all pending requests. This command causes Gx-Plus to send a JSER message to PCRF that includes the Juniper-Event-Type AVP (AVP code 2103) with a value of 3 indicating a discovery request. The PCRF then returns a JDER message to initiate discovery of all subscribers. When this discovery completes, all pending subscriber requests are cleared.

PCRF-Generated Discovery

The PCRF runs a discovery process in response to data loss, exhaustion of router resources, operator request, or router request. The JSDR message specifies the level of verbosity desired in the reply from Gx-Plus. The message also specifies whether the request is for data about a particular session or information similar to an SNMP Get-Bulk for all sessions. Gx-Plus returns a JSDA message that indicates complete success, limited success, or an error. In the event of success, the requested data is also returned.

Subscriber Accounting

When the PCRF returns a CCA-I message to the router, the message may contain thresholds for any of several usage statistics for a subscriber session or service session: Duration, input data, output data, or total data for the session. Upon receipt of a threshold, the router begins monitoring the subscriber's service session activity for that statistic. When the usage statistic reaches the threshold, it triggers the router to send a Gx-Plus usage notification message (CCR-U) to the PCRF. In response, the PCRF may send a CCA-U message to specify a new threshold, activate new services, or deactivate current services.

The PCRF can also send a CCR-U message that explicitly requests usage monitoring for statistics at different levels. The router can monitor usage at the subscriber level or at the service level. The Granted-Service-Unit AVP in the message specifies one or more of the following the statistics:

- CC-Input-Octets
- CC-Output-Octets
- CC-Total-Octets
- CC-Time

If any other statistics are specified, the router sends the PCRF a CCA message indicating that incorrect statistics were requested. When the specified threshold for a monitored statistic is reached, the router sends a CCR-U that contains the usage report for the statistics. In response, the PCRF sends another CCA-R with new thresholds or a request to activate or deactivate services.

Subscriber Usage Thresholds

Gx-Plus threshold monitoring enables the tracking of session statistics including the duration of session and the number of input bytes, output bytes, and total bytes allowed (granted) and used. Threshold monitoring involves the use of numerous AVPs.

- Rule-Install AVP—a grouped AVP that can consist of the following two AVPs:
 - Rule-Install-Name AVP—The name of the dynamic-profile to activate, corresponding to a service.

- Monitoring-Key AVP—(Optional) The name of the monitoring definition, which is part of the CCR/RAR messages, and indicates that Gx-Plus thresholds are enabled. The Monitoring-Key AVP must be unique within the context of the subscriber, but more than one of these keys can be included in the Rule-Install AVP, one per subscriber. For every Monitoring-Key AVP referenced in the Rule-Install AVP, there must be a corresponding Monitoring AVP.
- Monitoring AVP—The monitoring definition, consisting of the Monitoring-Key AVP and either the Granted-Service-Unit AVP or the Used-Service-Unit AVP:
 - Monitoring-Key AVP—The name of the monitoring definition.
 - Granted-Service-Unit AVP—A grouped AVP that includes the following session threshold values:
 - Duration AVP—Period of time in seconds allotted to the subscriber before having to ask for an extension.
 - Input-Bytes AVP—Number of input bytes allotted to the subscriber before having to ask for an extension. A value of zero indicates the threshold is turned off.
 - Output-Bytes AVP—Number of output bytes allotted to the subscriber before having to ask for an extension. A value of zero indicates the threshold is turned off.
 - Total-Bytes AVP—Number of input and output bytes in total allotted to the subscriber before having to ask for an extension.

The Granted-Service-Unit threshold values are somewhat analogous to a lease. In this case, if no threshold values are supplied, then the granted values or “lease” is effectively infinite. The absence of thresholds means no limits are placed on the values.

- Used-Service-Unit AVP—A grouped AVP that includes the following session threshold values, which are analogous to a kind of lease:
 - Duration AVP—Period of time in seconds that the service has been used.
 - Input-Bytes AVP—Number of input bytes used by the subscriber in this session.
 - Output-Bytes AVP—Number of output bytes used by the subscriber in this session.
 - Total-Bytes AVP—Number of input and output bytes in total used by the subscriber in this session.

No thresholds are enabled if the router acting as a PCEF receives a CCA or RAR message that contains one or more Rule-Install-AVPs, but no Monitoring-Key AVPs.

Consider the following example. The PCEF receives the listed AVPs in a CCA-I message. When the PCEF activates the svc-21-g service, the set of monitored thresholds, thresh-459 becomes active for the service. The instantiated service is granted 600 seconds, 1 billion input bytes, 1 billion output bytes, and a total of 2 billion bytes combined.

- Rule-Install AVP
 - Rule-Install-Name AVP = svc-21-g
 - Monitoring-Key AVP = thresh-459
- Monitoring AVP
 - Monitoring-Key AVP = thresh-459
 - Granted-Service-Unit AVP
 - Duration AVP = 600s
 - Input-Bytes AVP = 1,000,000,000
 - Output-Bytes AVP = 1,000,000,000
 - Total-Bytes AVP = 2,000,000,000

If the CCA-I includes the following AVPs and values, everything is the same as above except that no limits are placed on either input bytes or output bytes, just a limit on the total number of bytes. Omitting the Input-Bytes and Output-Bytes AVPs from the Granted-Service-Unit AVP has the same effect.

- Rule-Install AVP
 - Rule-Install-Name AVP = svc-21-g
 - Monitoring-Key AVP = thresh-459
- Monitoring AVP
 - Monitoring-Key AVP = thresh-459
 - Granted-Service-Unit AVP
 - Duration AVP = 600s
 - Input-Bytes AVP = 0
 - Output-Bytes AVP = 0
 - Total-Bytes AVP = 2,000,000,000

It does not matter which threshold is met first; the PCEF behaves the same.

1. It disables the complete set of monitored thresholds for the service. In the examples above, thresh-459 is disabled for service svc-21-g.
2. Authd sends a threshold report (CCR-U) to the PCRF that includes the Monitoring AVP with the current values for the thresholds; these make up the Used-Service-Unit AVP:

- Monitoring AVP
 - Monitoring-Key AVP = thresh-459
 - Used-Service-Unit AVP
 - Duration AVP = 600s
 - Input-Bytes AVP = 22,110,000
 - Output-Bytes AVP = 21,161,004
 - Total-Bytes AVP = 43,271,004

3. authd expects the PCRF to respond to the CCR-U with the Monitoring AVP, supplying new values for the thresholds. To use the lease analogy, the reply should extend the “lease” for the session; for example..

- Monitoring AVP
 - Monitoring-Key AVP = thresh-459
 - Granted-Service-Unit AVP
 - Duration AVP = 3600s
 - Input-Bytes AVP = 1,500,000,000
 - Output-Bytes AVP = 2,000,000,000
 - Total-Bytes AVP = 3,500,000,000

If the new Duration AVP supplied by the PCRF is low, it could result in a tight cycle of threshold hits, reports, and updates. Consequently the PCEF ensures that the threshold is of a reasonable duration by adding the new value from the PCRF to the current reported value; this becomes the new duration grant. Using the example above, the (current value + new value) = 600 + 3600 = 4200 seconds.

What happens if the PCRF fails to respond to the CCR-U? Rather than leave the thresholds disabled, the PCEF supplies the Monitoring AVP with a single new value, the duration:

- Monitoring AVP
 - Monitoring-Key AVP = thresh-459
 - Granted-Service-Unit AVP
 - Duration AVP = *current value + minimum-duration*

The router has default minimum values for all the threshold AVPs:

- Input-Bytes minimum - 1,000,000

- Output-Bytes minimum - 1,000,000
- Total-Bytes minimum - 1,000,000
- Duration minimum - 600

Using the example of 600 seconds for the current duration value, if the PCRF does not respond to the CCR-U, the new duration value becomes $600 + 600 = 1200$ seconds. There are no thresholds for the byte counts. When the new duration threshold is met, the PCEF generates another CCR-U threshold report for the PCRF.

Subscriber Audit

The PCRF can send a reauthorization request (RAR message) to Gx-Plus at any time to determine whether a particular subscriber is still logged in. You can also manually trigger the PCRF to do so by issuing the clear network-access aaa gx-plus replay command.

The Session-Id AVP identifies the subscriber session. Gx-Plus returns an RAA message to provide status on the subscriber session. When the session is still up (found in the session database) the Result-Code AVP value in the RAA message is DIAMETER_SUCCESS (2001). When the session is not found, the Result-Code value is DIAMETER_UNKNOWN_SESSION_ID (5002). A Result-Code value of DIAMETER_UNABLE_TO_DELIVER (3002) indicates that Gx-Plus is not configured.

Starting in Junos OS Release 17.4R1, the router updates monitored statistics when they are received in the RAR from the PCRF. When Gx-Plus sends an RAA message after receiving an RAR message requesting service activation or deactivation, it also sends a CCR-U message to the PCRF with updated statistics.

Subscriber Logout

When the client application sends a subscriber logout notice to AAA, Gx-Plus sends a CCR-T message to notify the PCRF that the provisioned subscriber session is being terminated. The PCRF returns a CCA-T message that includes the Result-Code AVP. If the Result-Code value is DIAMETER_SUCCESS, Gx-Plus notifies AAA, and AAA notifies the application that the logout is complete. If Gx-Plus does not receive a CCA-T message, or if the Result-Code AVP has any other value or is missing, then the termination request is retried until the CCA-T message is returned with DIAMETER_SUCCESS.

SEE ALSO

[Default Subscriber Service Overview | 1091](#)

[Configuring a Default Subscriber Service | 1092](#)

Configuring Gx-Plus

You can configure the Gx-Plus client application to work with a PCRF policy manager residing on a server. The PCRF is a centralized policy decision point that deploys business rules to allocate broadband network resources and manage subscribers and services. AAA on the router (acting as the PCEF) uses Gx-Plus to request service provisioning from the PCRF.



NOTE: Contact the Juniper Networks Technical Assistance Center (JTAC) for information on supported PCRFs.

To configure Gx-Plus:

1. Configure the Gx-Plus partition.
See ["Configuring the Gx-Plus Partition" on page 1794](#).
2. Configure Gx-Plus global attributes: the number of outstanding requests permitted and the inclusion of IPv6 subscribers.
See ["Configuring Gx-Plus Global Attributes" on page 1795](#).
3. Configure Gx-Plus provisioning for subscribers.
See ["Provisioning Subscribers with Gx-Plus" on page 1796](#).
4. (Optional) Override PCRF control of a subscriber session to correct services or troubleshoot a problem.
See ["Disabling PCRF Control of a Subscriber Session" on page 1797](#).
5. (Optional) Configure Gx-Plus event tracing as part of general authentication service tracing operations.
See [Tracing General Authentication Service Processes](#).

Configuring the Gx-Plus Partition

Gx-Plus works within a specific logical system: routing instance context, called a partition.



NOTE: Currently, only a single partition is supported; you must configure it within the default logical system: routing instance context.

Before you configure the Gx-Plus partition, perform the following task:

- Configure the Diameter instance at the [edit diameter] hierarchy level. See ["Configuring Diameter" on page 1762](#).

Configuration for the Gx-Plus partition consists of naming the partition and then associating a Diameter instance, the PCRF hostname, and the PCRF realm with the partition.

To configure the Gx-Plus partition:

1. Create the partition or specify the name of an existing partition.

```
[edit access gx-plus]
user@host# set partition partition-name
```

2. Specify the Diameter instance for the Gx-Plus partition.



NOTE: Currently, only the default Diameter instance, `master`, is supported.

```
[edit access gx-plus partition partition-name]
user@host# set diameter-instance instance-name
```

3. (Optional) Configure the destination host for the Gx-Plus partition.

```
[edit access gx-plus partition partition-name]
user@host# set destination-host hostname
```

4. Configure the destination realm for the Gx-Plus partition.

```
[edit access gx-plus partition partition-name]
user@host# set destination-realm realm
```

The following example shows a Gx-Plus partition configuration.

```
gx-plus {
  partition partition1 {
    diameter-instance master;
    destination-host pcrf1;
    destination-realm generic.example.com;
  }
}
```

Configuring Gx-Plus Global Attributes

You can configure attributes that apply to all Gx-Plus partitions globally.

When a request from Gx-Plus to the PCRF is not answered or is improperly answered, Gx-Plus keeps retrying the request until it receives an appropriate answer. If the number of requests grows too large,

the PCRF can become overloaded and messages can be lost. To reduce this risk, you can set a limit on the number of outstanding requests to the PCRF that Gx-Plus can retry.

By default, Gx-Plus does not include IPv6 subscribers in Gx-Plus provisioning requests to the PCRF. Instead, Gx-Plus only establishes sessions that correspond to IPv4 DHCP sessions on dual-stack IPv6/IPv4 or IPv4-only subscriber interfaces. You must explicitly configure Gx-Plus to include IPv6 information. When you do so, Gx-Plus can establish sessions that correspond to DHCPv6 sessions on IPv6-enabled subscriber interfaces and on dual-stack IPv6/IPv4-enabled interfaces.

To configure Gx-Plus global attributes:

1. (Optional) Set a limit on the number of outstanding requests.

```
[edit access gx-plus global]
user@host# set max-outstanding-requests number
```

2. (Optional) Include IPv6 subscribers in provisioning requests.

```
[edit access gx-plus global]
user@host# set include-ipv6
```

For example to limit the number of outstanding requests to 30 and to include IPv6 subscribers:

```
[edit access gx-plus global]
user@host# set max-outstanding-requests 30
user@host# set include-ipv6
```

Provisioning Subscribers with Gx-Plus

You can configure AAA to use Gx-Plus to request provisioning from a PCRF to instantiate services for an authenticated subscriber.

Before you configure Gx-Plus provisioning for subscribers, perform the following task:

- Create the subscriber access profile at the [edit access profile] hierarchy level.

To configure Gx-Plus provisioning:

- Specify `gx-plus` as the provisioning method in the profile.

```
[edit access profile profile-name]
user@host# set provisioning-order gx-plus
```

Disabling PCRF Control of a Subscriber Session

When a subscriber has been provisioned with Gx-Plus, services for that subscriber can be activated and deactivated only by the PCRF. Accordingly, AAA rejects any RADIUS CoA requests for subscribers provisioned by Gx-Plus. Similarly, CLI-based service activation and deactivation do not work while a subscriber is remotely provisioned.

Network administrators without PCRF access or authority may need to override PCRF control on a particular subscriber session to troubleshoot the session or correct the subscriber services. You can disable PCRF control by issuing the `request network-access aaa subscriber set session-id` command. In response, the router sends a termination notice to the PCRF, but does not actually log out the subscriber.

When you have confirmed that provisioning is disabled, you can then activate or deactivate subscriber services for that session with the `request network-access aaa subscriber add session-id` and `request network-access aaa subscriber delete session-id` commands, respectively. These commands fail if provisioning is still enabled.

Another consequence of disabling provisioning for a subscriber session is that RADIUS change of authorization (CoA) messages can modify the session.

Before you begin, determine or verify the ID for the session by displaying the session IDs of all current subscribers with the `show subscribers detail` or `show network-access aaa subscribers` command.

To disable control by the PCRF over a subscriber session:

1. Disable provisioning for the specified subscriber session ID.

```
user@host> request network-access aaa subscriber set session-id subscriber-session-id
provisioning-state none
```

2. (Optional) Verify that provisioning is disabled for the session.

```
user@host> show network-access aaa subscribers session-id subscriber-session-id detail
```

For example, to disable provisioning for subscriber larry:

```
user@host> show network-access aaa subscribers
Username      Logical system/Routing instance  Client type  Session-ID
...
larry         default:default                  dhcp         55
...
user@host> request network-access aaa subscriber set session-id 55 provisioning-state none
```

```
user@host> show network-access aaa subscribers session-id 55 detail
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 55
Accounting Session ID: 55
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Accounting State: Acc-Start-Send
Provisioning-type: none
Service name: basic-service
Service State: SvcActive
Session ID: 56
Session uptime: 00:01:45
```

SEE ALSO

| *Local and Remote Service Activation and Deactivation Using the CLI*

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the CCR-I message includes the Subscription-Id AVP (AVP code 443) with the Subscription-Id-Type AVP set to 4 and Subscription-Id-Data AVP set to reserved.
17.4R1	Starting in Junos OS Release 17.4R1, the router updates monitored statistics when they are received in the RAR from the PCRF. When Gx-Plus sends an RAA message after receiving an RAR message requesting service activation or deactivation, it also sends a CCR-U message to the PCRF with updated statistics.

RELATED DOCUMENTATION

| [Diameter Base Protocol](#) | 1727

[3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1799](#)

[NASREQ for Authentication and Authorization | 1852](#)

[JSRC for Subscriber Provisioning and Accounting | 1856](#)

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

IN THIS SECTION

- [3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1799](#)
- [Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1802](#)
- [Understanding Gx Interactions Between the Router and the PCRF | 1807](#)
- [Understanding Gy Interactions Between the Router and the OCS | 1819](#)
- [Gy File Backup Overview | 1826](#)
- [Understanding Interactions Between the PCRF, PCEF, and OCS | 1828](#)
- [Understanding Upstream and Downstream Messages for the PCRF | 1833](#)
- [Configuring the OCS Partition | 1837](#)
- [Configuring the PCRF Partition | 1844](#)
- [Configuring OCS Global Parameters | 1851](#)

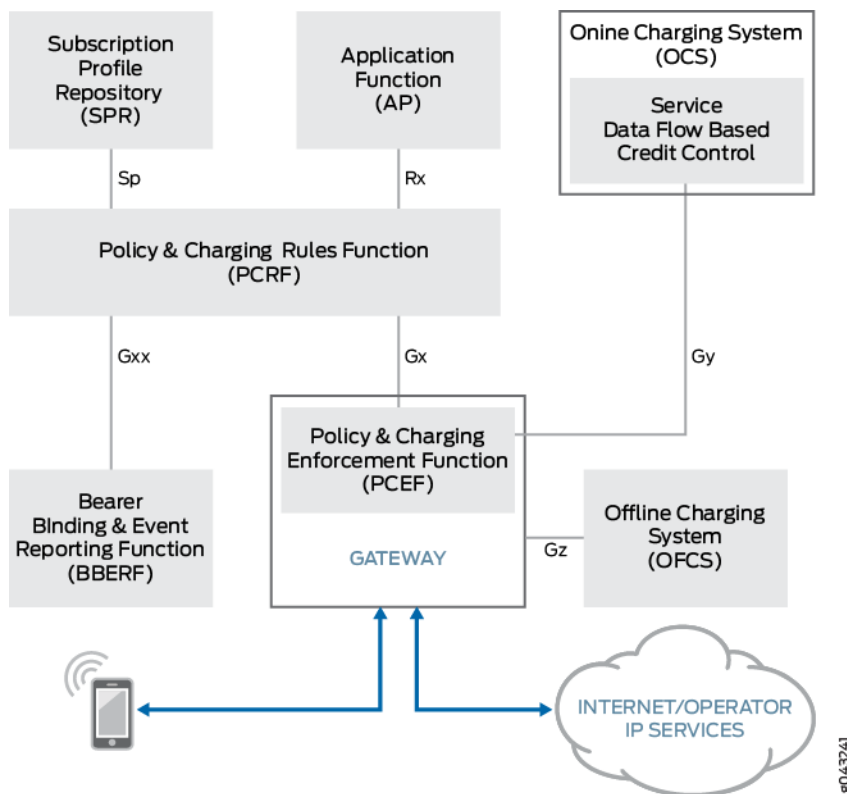
3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

IN THIS SECTION

- [Benefits of 3GPP Policy and Charging Control Architecture | 1801](#)

The 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) provides the unification of wireline provisioning and accounting for customers. [Figure 51 on page 1800](#) shows the components of an overall 3GPP PCC architecture.

Figure 51: 3GPP PCC Architecture Overview



The four major components of the PCC architecture are:

- **Policy and Charging Rules Function (PCRF)**—A centralized policy decision point that deploys business policy and charging rules to allocate broadband network resources and manages flow-based charges for subscribers and services. PCRF pushes the rules down to the Policy and Charging Enforcement Function (PCEF) using the 3GPP Gx protocol and online policy interface.
- **Policy and Charging Enforcement Function (PCEF)**—A function that provides user traffic handling and QoS at the gateway, provides service data flow detection, and applies the rules received from the PCRF. PCEF optionally interacts with the Online Charging Function (OCF) within the Online Charging System (OCS) using the 3GPP Gy protocol to retrieve policy and charging authorization for quotas and credit control.
- **Online Charging System (OCS)**—The component responsible for interacting with the PCEF. The PCEF optionally reports usage and receives additional authorizations from the OCS using the 3GPP Gy protocol. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.
- **Offline Charging System (OFCS)**—A process where charging information for network resource usage is collected concurrently with that resource usage. If credit-based authorization is not required, the PCEF applies policies and report usage to the OFCS using the 3GPP Gz protocol. You can also use the OCS as the primary accounting destination and use the OFCS as a backup.

Table 79 on page 1801 lists the functionality differences between PCRF and PCEF.

Table 79: Functionality Comparison Between PCRF and PCEF

Functionality	PCRF	PCEF
Charging policing implementation	Involved at different levels; aggregates information inside the hosting network and is considered part of the PCC architecture.	Involved at different levels; located at the gateway.
Functions included	Includes mainly policy control decision and flow-based control functions.	Includes policy enforcement and flow-based charging functions.
Predefined PCC rules	Activation or deactivation of predefined PCC rules can only be done by the PCRF.	Preconfigured by the PCEF.
Online and offline charging interactions	Not supported	Supported

The three other components that make up the PCC architecture in Figure 51 on page 1800 are:

- **Application Function (AF)**—The Application Function interacts with applications or services that require dynamic PCC. The Application Function extracts session information from the application signalling and provides application session-related information to the PCRF using the Rx protocol.
- **Subscription Profile Repository (SPR)**—SPR contains subscriber and subscription information on a per-packet data network (PDN) basis. The Sp protocol enables the PCRF to request subscription information related to a subscriber's service or session.
- **Bearer Binding and Event Reporting Function (BBERF)**—The PCC rule needs to be mapped to a particular IP bearer to ensure the packets receive the appropriate QoS treatment. The association between a PCC rule and a bearer is referred to as *bearer binding*. The BBERF location depends on the access technology. For 3GPP, the BBERF is located in the serving gateway and uses the Gxx protocol.

Benefits of 3GPP Policy and Charging Control Architecture

- Provides a unified framework for wireline subscriber provisioning and accounting.


Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

IN THIS SECTION

- Wireline Access Environment | 1803
- Junos OS Environment | 1805

The Policy and Charging Enforcement Function (PCEF) is one of four major components of the 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) architecture in [Figure 52 on page 1802](#).

Figure 52: 3GPP PCC Architecture Overview



PCEF provides user traffic handling and quality of service (QoS) at the gateway, provides service data flow detection, and applies the rules received from the Policy Control and Charging Rules Function (PCRF). 3GPP defines Gx as the online policy protocol between the PCRF and the PCEF to provide control over policy and flow-based charges for subscribers. The PCRF is a centralized policy decision

point that deploys business policy rules to allocate broadband network resources and manages flow-based charges for subscribers and services. Optionally, the PCEF interacts with the Online Charging System (OCS) using the 3GPP Gy protocol to retrieve policy and charging authorization for quotas and credit control.

PCEF provides support for the following environments:

Wireline Access Environment

For mobile subscribers, the user equipment requests services; for broadband wireline subscribers, the PCRF requests services. In the wireline environment, PCRF functions as the service requester, and the PCEF functions as the service receiver and enforcer.

Adapting the PCC model in a wireline environment provides these benefits:

- Convenience
- Advanced technology
- Already implemented by the wireless branch of the carrier that often provides a much bigger business than the wireline branch

The PCRF controls the PCEF by pushing charging rules. Charging rules are reused as service (policy) rules to push policies. Charging rules may also have an associated rating group, or charging key. As a result, the PCEF configuration must define charging rules and mapping between credit control services (cc-services) and rating groups.

In many instances, both OCS and Offline Charging System (OFCS) 3GPP accounting services require Mobile Station International Subscriber Directory Number (MSISDN) be used for subscriber identification. The MSISDN is passed as the subscription ID. While each mobile user equipment device has an associated MSISDN, this information is not available for wireline subscribers. To enable the PCRF to dynamically pass subscription-ID parameters, and support a variety of authentication, authorization, and provisioning configuration, the Juniper attribute-value pairs (AVPs) in [Table 80 on page 1804](#) have been allocated from the Juniper Vendor-ID space (2636) vendor-specific attribute (VSA).



NOTE: If no dynamic-subscription ID is received, then neither OCS or OFCS communications are initiated.

Table 80: Allocated Juniper AVPs

AVP Name	Vendor-ID	AVP Type	Diameter Type	Diameter Flag
Juniper-Dyn-Subscription-Indicator	2636	10001	Enum	V
Juniper-Dyn-Subscription-Id	2636	10002	Grouped	VM
Juniper-Dyn-Subscription-Id-Type	2636	10003	Integer32	VM
Juniper-Dyn-Subscription-Id-Data	2636	10004	UTF8String	VM

The client system (router) sends the Juniper-Dyn-Subscription-Id-Indicator AVP to indicate support of the dynamic assignment of the subscription ID. The Juniper-Dyn-Subscription-Id-Indicator attribute has two values:

- DYN_SUBSCRIPTION_NOT_SUPPORTED (0)
- DYN_SUBSCRIPTION_SUPPORTED (1)

The server then sends the Juniper-Dyn-Subscription-Id AVP to the client that indicated support. This is a grouped AVP that contains the values to be sent as Subscription-Id-Type and Subscription-Id-Data.

**NOTE:**

- The PCRF server may use standard Subscription-Id AVP to communicate the dynamic-subscription ID to the router.
- If both Juniper-Dyn-Subscription-Id and Subscription-Id are sent by the PCRF, the Subscription-Id value is used.

In many cases, wireline subscribers support only one IP family, which is required information for both AAA service and PCRF. To indicate this information, the Juniper-Network-Family-Indicator AVP has been allocated from the Juniper Vendor-ID space (2636) VSA in [Table 81 on page 1805](#).

Table 81: Family Indicator AVP

AVP Name	Vendor-ID	AVP Type	Diameter Type	Diameter Flag
Juniper-Network-Family-Indicator	2636	10010	Enum	V

The client system (router) sends the Juniper-Network-Family-Indicator AVP to indicate which network families are associated with the service request and supported by the subscriber. When you configure the Juniper-Network-Family-Indicator AVP to indicate the associated network family, the system sends the information to the PCRF. The Juniper-Network-Family-Indicator attribute has four values:

- UNSPECIFIED (0)
- IPV4_FAMILY (1)
- IPV6_FAMILY (2)
- IPV4_IPV6_FAMILY (3)

Wireline customers often control user services solely through the PCRF and use the OCS as a convenient real-time usage monitoring mechanism rather than as an enforcement unit. To decrease the number of possible erroneous OCS configurations, include the `force-continue` statement at the `[edit access ocs partition partition-name]` hierarchy level to force the broadband PCEF (BPCEF) to limit the impact of negative responses from the OCS and quota expirations, and to prevent sending OCS notifications for affected rating-groups. Whenever the PCEF receives a negative response to any reported group, it stops reporting this group to the OCS.

Junos OS Environment

There are three categories of dynamic-profiles within the Junos OS environment:

- client-dynamic-profiles
- cos-service-dynamic-profiles
- firewall-service-dynamic-profiles

Client-dynamic-profiles and cos-service-dynamic-profiles define bandwidth and other characteristics of the services provided to a subscriber; the firewall-service-profiles perform filtering and usage counting. For all of the dynamic-profiles' categories, the service-dynamic-profile name is used as the value of a Charging-Rule-Name AVP.

When the service-dynamic-profile has no variables, or when defaults provided in service-dynamic-profile definition are requested, no additional elements are required. To provide custom values for a service-dynamic-profile, use the Charging-Rule-Definition AVP with additional VSAs.

The PCRF uses existing Juniper-Substitution VSAs (Vendor-ID 2636 and Type 2024) to supply attributes as a name-value pairs. The PCRF may also include parameters as positional notation for part of the rule name. The Redirect-Information AVP (Vendor-ID 10415 and Type 1085) supplies a value for the Redirect-URL parameter.

For every possible service-dynamic-profile parameter name requested by customers, a new Juniper-Parameter VSA is defined. [Table 82 on page 1806](#) describes the initial set of fixed Juniper-Parameter VSAs.

Table 82: Initial Set of Fixed Juniper-Parameter VSAs

Parameter	VSA Name	Vendor-ID	Type	Diameter Type
Cos-Tcp	Juniper-Param-Cos-Tcp	2636	10005	UTF8String
V4-Firewall-Input-Filter	Juniper-Param-V4-Firewall-Input-Filter	2636	10006	UTF8String
V4-Firewall-Output-Filter	Juniper-Param-V4-Firewall-Output-Filter	2636	10007	UTF8String
V6-Firewall-Input-Filter	Juniper-Param-V6-Firewall-Input-Filter	2636	10008	UTF8String
V6-Firewall-Output-Filter	Juniper-Param-V6-Firewall-Output-Filter	2636	10009	UTF8String

If parameters or the Service-Identifier and Rating-Group are required to be indicated by the PCRF, the Charging-Rule-Definition AVP is used; otherwise, the Charging-Rule-Name AVP is used.

```
Charging-Rule-Definition ::= < AVP Header: 1003 >
    { Charging-Rule-Name }
    [ Service-Identifier ]
    [ Rating-Group ]
    [ Online ]
    [ Precedence ]
    [ Juniper-Param-VSA ]
    [ AVPs ] - standard AVPs used as parameters
```

For instances when there is a Service-Identifier and Rating-Group combination, or when only the Service-Identifier or only the Rating-Group is specified, the combination must be unique among the rules installed for a subscriber. You configure the service-context-id on the router.

Understanding Gx Interactions Between the Router and the PCRF

IN THIS SECTION

- [Subscriber Login | 1807](#)
- [Subscriber Login Error Recovery | 1812](#)
- [Subscriber Update | 1814](#)
- [Subscriber Logout | 1816](#)
- [Subscriber Disconnect | 1818](#)
- [Connectivity Fault Recovery | 1819](#)

The sequences of Diameter messages are exchanged by means of the 3rd Generation Partnership Project (3GPP) Gx protocol between the Policy Control and Rules Charging Function (PCRF) and the router acting as a Policy and Charging Enforcement Function (PCEF).

Starting in Junos OS Release 17.3R1, support for additional OCS and PCRF features are added using Gy and Gx protocols. The new statements:

- `accept-sdr` is added for PCRF partition at the hierarchy level `[edit access pcrf partition partition-name]`.
- `alternative-partition-name` is added for OCS partition at the hierarchy level `[edit access ocs partition partition-name]`.

They interact to perform the following subscriber access tasks:

Subscriber Login

The router sends a Diameter CCR request containing a fixed set of required information to a policy manager (PCRF) and receives a CCA response containing policies and other information. Gx provisioning is enabled for subscribers when you include the `provisioning-order pcrf` statement at the `[edit access profile profile-name]` hierarchy level. When an application requests AAA to activate the subscriber's session, the router sends a CCR-GX-I (where I represents INITIAL_REQUEST) message to the PCRF to request a fix set of provisioning information for the subscriber session, and receives a CCA-GX-I response message containing policies and other information, including the Result-Code AVP (AVP code 268).

When you configure the `provisioning-order` statement in the access profile, the broadband PCEF (BPCEF) module sends a provisioning request to the PCRF during the client activation. The following examples show a CCR-GX-I and CCA-GX-I packet exchange:

CCR-GX-I Packet Example

```
CCR-GX-I ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <configurable-string> }
{ Origin-Realm:                  <configurable-string> }
{ Destination-Realm:             <configurable-string> }
{ CC-Request-Type:               INITIAL_REQUEST(1) }
{ CC-Request-Number:             0 }
{ Subscription-Id:
    { Subscription-Id-Type:       <configurable-integer> }
    { Subscription-Id-Data:       <configurable-string> }
}
}
[ Destination-Host:              <configurable-string> ] -- if configured
[ Origin-State-Id:                <u32> ] -- if configured to send
[ Framed-IP-Address:              <ipv4-address-in-radius-encoding> ] -- if available
[ Framed-IPv6-Prefix:             <ipv6-prefix-in-radius-encoding> ] -- if available
{ IP-CAN-Type:                    <configurable-integer> }
{ Online:                         ENABLE_ONLINE (1) }
[ User-Name:                      <string> ]
[ NAS-Port-Id:                    <string> ] -- if included by config
[ Juniper-Virtual-Router:         <virtual-router-name> ] -- if included by config
[ Event-Timestamp:                <timestamp> ] -- login timestamp, if included by config
{ Juniper-Dyn-Subscription-Indicator: DYN_SUBSCRIPTION_SUPPORTED(1) }
{ Juniper-Network-Family-Indicator: <subscriber-family> }
```



NOTE: The T (potentially retransmitted message) bit recalculates when the CCR-GX-I is resent. This flag is set after a link failover procedure to remove duplicate requests.

CCA-GX-I Packet Example

```
CCA-GX-I ::= <Diameter Header: 272, PXY, 16777238>
{ <Session-Id> }
{ Result-Code:                   <integer> }
{ Auth-Application-Id:           16777238 }
{ Origin-Host:                   <string> } -- should match destination-host if configured
{ Origin-Realm:                   <string> } -- should match destination-realm
{ Result-Code:                   <integer> }
```

```

{ CC-Request-Type:          INITIAL_REQUEST(1) }
{ CC-Request-Number:        0 }
[ Juniper-Dyn-Subscription-Id:
  {Juniper-Dyn-Subscription-Id-Type:    <value-to-be-used-for-ocs-interactions> }
  {Juniper-Dyn-Subscription-Id-Data:    <value-to-be-used-for-ocs-interactions> }
]
*[ Supported-Features ]          -- ignored
[ Origin-State-Id:              <u32> ] -- Indicates restart PCRF side
*[ Downstream data units ]

```



NOTE: If no rule-install AVP is defined in the CCA-GX-I, then the default rule is installed. All event triggers, including those not yet defined, are acceptable. However, only a few event triggers actually generate events when implemented.

The PCRF returns a CCA-GX-I message that includes the Result-Code AVP (AVP code 268) that maps to the result categories listed in [Table 83 on page 1809](#).

Table 83: Result-Code-AVP Categories

Result-Code-AVP Value	Result Category
SUCCESS(2001), LIMITED_SUCCSS(2002), and valid message	Grant
AUTHENTICATION_REJECTED(4001), UNKNOWN_SESSION_ID(5002), AUTHORIZATION_REJECTED(5003), and USER_UNKNOWN(5030)	Deny
UNABLE_TO_DELIVER(3002), REALM_NOT_SERVED(3003), TOO_BUSY(3004), LOOP_DETECTED(3005), and REDIRECT_INDICATION(3006)	Failure
All other Diameter Permanent-failure Result-Code AVPs greater than and equal to 5000, and all Diameter protocol error Result-Code AVPs greater than and equal to 3000 and less than 4000	Permanent-failure
Other Result-Code AVPs for invalid message or no-response	Failure

As shown in [Table 84 on page 1810](#), the CCA-GX-I response processing depends on three factors:

- Whether the local decision timeout has expired
- The setting of the local decision

- The result category

Table 84 on page 1810 also contains PCRF local decision timeout expiration actions.

Table 84: CCA-GX-I Response Processing

PCRF Local Decision Timeout	PCRF Local Decision	Result Category	Action
Not-expired	–	Grant	Clear the local decision timer, apply rules from the CCA-GX-I, notify the Online Charging System (OCS), and then acknowledge subscriber activation.
Not-expired	–	Deny	Clear the local decision timer and fail subscriber activation.
Not-expired	–	Failure	Retry the CCA-GX-I until the local decision time outs.
Not-expired	Grant	Permanent-failure	Clear the local decision timer, apply the default rule, acknowledge subscriber activation, and then keep retrying the CCA-GX-I.
Not-expired	Deny	Permanent-failure	Fail the subscriber activation and initiate the subscriber logout process.
On-expiration	Grant	–	Apply the default rule, keep retrying the CCA-GX-I indefinitely, and acknowledge subscriber activation.
On-expiration	Deny	–	Fail the subscriber activation and initiate the subscriber logout process.
Expired	Grant	Grant	If the CCA-GX-I contains rules, remove the default rules and install the received rules, and then notify the OCS and acknowledge subscriber activation.

Table 84: CCA-GX-I Response Processing (Continued)

PCRF Local Decision Timeout	PCRF Local Decision	Result Category	Action
Expired	Grant	Deny	Log out the client.
Expired	Grant	Failure	Keep retrying the CCA-GX-I indefinitely.
Expired	Grant	Permanent-failure	Take a long pause and then restart retrying the CCA-GX-I.
Expired	Deny	Deny	If subscriber still logging out, ignore subscriber; otherwise, no action required.

A subscriber login initiates the following sequence of events:

1. A client application—such as DHCP, PPP, or static subscriber sessions—requests AAA to authenticate the subscriber.
2. Authentication begins if the subscriber access profile specifies RADIUS authentication. Login continues when the authentication is successful. Login fails when the authentication-order statement in the profile does not specify RADIUS authentication or no authentication. Login also fails when authentication fails.
3. Default services are activated for the subscriber. Any services that the authentication server includes in the authentication grant are activated. Additionally, a default service may have been configured for the client application.
4. If the subscriber access profile specifies Gx provisioning, the router initiates the Gx message exchange by sending a CCR-GX-I message to the PCRF. The router waits for the PCRF to respond with a CCA-GX-I message within a non-configurable timeout period.

When the PCRF responds within the timeout period and includes the Charging-Rule-Install AVP in the CCA-GX-I message, subscriber login is delayed while the router deactivates any default services and attempts to activate the specified services.

- If all the specified services are activated, then the login completes.
- If any of the services cannot be activated, the router sends the PCRF a CCR-GX-U (where U represents UPDATE_REQUEST) message with the status of the services (a rule report). The PCRF responds to this message with a CCA-GX-U that can contain a new set of services for activation.

- The router ignores any default services, even if the CCA-GX-I message does not include any services. In this circumstance, no services are activated.

If the PCRF does not return a CCA-GX-I within the timeout period, subscriber login completes.

- The router searches first for services returned from the authentication server and activates any it finds. If no such services are found, then the router activates any locally configured default services. Subscriber login completes when default service activation is successful, but fails when any default service fails to activate. Because default services are not required to be present, login also completes when no default services are found.
- If login completes (with or without a default service), the router periodically resends the CCR-GX-I message to the PCRF. If the PCRF subsequently returns a CCA-GX-I, the router deactivates the default service, if any, and then activates any services included in the CCA-GX-I. If the message does not include any services, then no services are activated, not even a default service.
- If any of the services contained in the CCA-GX-I cannot be activated, the router sends the PCRF a CCR-GX-U message with the status of the services (a rule report). The PCRF responds to this message with a CCA-GX-U that can contain a new set of services for activation.

Subscriber Login Error Recovery

Starting in Junos Release 20.1R1, you can configure the router to recover from certain PCRF server errors by reinitializing the subscriber session to resync the router and PCRF server states. Some PCRF servers might not properly handle a situation where the CCA-GX-I messages that it sent to the router are lost. When the router retries sending the CCR-GX-I to the PCRF, the server is out of sync with the router because it has already sent a reply and is not aware that the router did not receive the message. This mismatch in state can lead to either of the following errors:

- The PCRF server responds to the retry with a CCA-GX-I that contains the Diameter Result Code AVP (Code 268) with a value of 5012 (DIAMETER UNABLE TO COMPLY). This is considered a permanent failure ([Table 83 on page 1809](#)).
- The PCRF server sends a RAR. The server expects the session to be active because it sent the CCA-GX-I to the router and is unaware that the message was not received. The server might send any of the following RAR messages:
 - RAR-GX-D to disconnect the session because it considers the session to be bad
 - RAR-GX-A to read information about the bad session
 - RAR-GX-U to update the session because it considers the session to be operating normally.

You can use the PCRF local-decision configuration to reinitialize the subscriber session in response to either or both of those errors.

- Include the `reinit-on-failure` option for the permanent failure error.

- Include `reinit-on-rar` option for the RAR error.



NOTE: The reinitialization operation has these additional configuration requirements:

- You must configure the local decision grant option.
- You must configure the router to use an extended session ID so that it can maintain state for the original session and the new one tied to the same login event. To do so, configure the PCRF `use-session-stamp` option.

The reinitialization operation consists of the following steps in both cases:

1. The router sends a session termination request, CCR-GX-T, to the PCRF to terminate the session. This is done in an attempt to get the router and PCRF server to have the same state for this session.
2. The router waits a reinitialization timeout period to receive a CCA-GX-T. You can use the `reinit-timeout` option to specify a period different than the default.
3. If the router either receives a CCA-GX-T within the timeout period or a CCA-GX-T does not arrive before the timeout expires, then the router sends a CCR-GX-I to the PCRF with a new, extended session ID. The extended session ID is conveyed in the Diameter Session-ID AVP (AVP code 263).

The router forms the extended session ID by appending a session stamp that consists of the UTC time when the router creates the CCR-GX-I. For example, consider the following Diameter Session-Id AVP. The session ID is 23 and `use-session-stamp` is not configured:

```
test-host1;0000000000;0000000023;
```

With `use-session-stamp` configured, the session timestamp is appended to the AVP value:

```
test-host1;0000000000;0000000023;1557788595;
```

[Table 85 on page 1813](#) provides details about how the router reacts to these errors based on the current local PCRF state.

Table 85: Router Actions Based on Local PCRF State

Local State	Action When PCRF Error Occurs
-------------	-------------------------------

local-active— Subscriber is active with default services.	<p>The router does the following:</p> <ul style="list-style-type: none"> • Transitions to the local-reinit state. • Sends a CCR-GX-T to the PCRF. • Starts the local-decision reinitialization timer and waits for the CCA-GX-T reply from the PCRF.
local-grant—Default service provisioning is in progress.	<p>When the default provisioning completes, the router does the following:</p> <ul style="list-style-type: none"> • Transitions to the local-reinit state. • Sends a CCR-GX-T to the PCRF. • Starts the local-decision reinitialization timer and waits for the CCA-GX-T reply from the PCRF.
started—The local- decision timer is still running.	<p>The router does the following when no default services are configured:</p> <ul style="list-style-type: none"> • Transitions to the local-reinit state. • Sends a CCR-GX-T to the PCRF. • Starts the local-decision reinitialization timer and waits for the CCA-GX-T reply from the PCRF. <p>The router does the following when default services are configured:</p> <ul style="list-style-type: none"> • Transitions to the local-reinit-early state. • Start provisioning the default services. <p>When the default provisioning completes, the router does the following:</p> <ul style="list-style-type: none"> • Transitions to the local-reinit state. • Sends a CCR-GX-T to the PCRF. • Starts the local-decision reinitialization timer and waits for the CCA-GX-T reply from the PCRF.

Subscriber Update

Whenever a trigger event occurs on the router, an update request is sent to the PCRF. The following examples show a CCR-GX-U (where U represents UPDATE_REQUEST) and CCA-GX-U packet exchange:

CCR-GX-U Packet Example

```
CCR-GX-U ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <configurable-string> }
{ Origin-Realm:                  <configurable-string> }
{ Destination-Realm:             <configurable-string> }
{ CC-Request-Type:               UPDATE_REQUEST(2) }
{ CC-Request-Number:             <u32> }
[ Destination-Host:              <configurable-string> ] -- if configured
[ Origin-State-Id:               <u32> ] -- if configured to send
*[ Upstream data units ]
```



NOTE: The T bit recalculates when the CCR-GX-U is resent.

CCA-GX-U Packet Example

```
CCA-GX-U ::= <Diameter Header: 272, PXY, 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <string> } -- should match destination-host if configured
{ Origin-Realm:                  <string> } -- should match destination-realm
{ Result-Code:                   <integer> }
{ CC-Request-Type:               UPDATE_REQUEST(2) }
{ CC-Request-Number:             <u32> }
[ Origin-State-Id:               <u32> ] -- Indicates PCRF restart
*[ Downstream data units ]
```

The PCRF returns a CCA-GX-U message that includes the Result-Code AVP (AVP code 268) that maps to the result categories listed in [Table 86 on page 1815](#).

Table 86: Result-Code-AVP Categories

Result-Code-AVP Value	Result Category
SUCCESS(2001), LIMITED_SUCCSS(2002), and valid message	Success

Table 86: Result-Code-AVP Categories (Continued)

Result-Code-AVP Value	Result Category
UNABLE_TO_DELIVER(3002), REALM_NOT_SERVED(3003), TOO_BUSY(3004), LOOP_DETECTED(3005), and REDIRECT_INDICATION(3006)	Failure
All other Diameter Permanent-failure Result-Code AVPs greater than and equal to 5000, and all Diameter protocol error Result-Code AVPs greater than and equal to 3000 and less than 4000	Success
Other Result-Code AVPs for invalid message or no-response	Failure

Subscriber Logout

When the client application sends a subscriber logout notice to AAA, Gx sends a CCR-GX-T (where T represents TERMINATION_REQUEST) message to notify the PCRF that the provisioned subscriber session is being terminated.

Whenever a trigger event occurs on the router, a terminate request is sent to the PCRF. The following examples show a CCR-GX-T and CCA-GX-T packet exchange:

CCR-GX-T Packet Example

```
CCR-GX-T ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <configurable-string> }
{ Origin-Realm:                  <configurable-string> }
{ Destination-Realm:             <configurable-string> }
{ CC-Request-Type:               TERMINATION_REQUEST(3) }
{ CC-Request-Number:             <u32> }
[ Destination-Host:               <configurable-string> ] -- if configured
{ Termination-Cause:             DIAMETER_LOGOUT(1) }
[ Origin-State-Id:               <u32> ] -- if configured to send
*[ Upstream data units ]
```



NOTE: The T bit recalculates when the CCR-GX-T is resent.

CCA-GX-T Packet Example

```
CCA-GX-T ::= <Diameter Header: 272, PXY, 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <string> } -- should match destination-host if configured
{ Origin-Realm:                  <string> } -- should match destination-realm
{ Result-Code:                   <integer> }
{ CC-Request-Type:               TERMINATION_REQUEST(3) }
{ CC-Request-Number:             <u32> }
[ Origin-State-Id:               <u32> ] -- Indicates PCRF restart
*[ Downstream data units ]
```

The PCRF returns a CCA-GX-T message that includes the Result-Code AVP (AVP code 268) that maps to the result categories listed in [Table 87 on page 1817](#).

Table 87: Result-Code-AVP Categories

Result-Code-AVP Value	Result Category
SUCCESS(2001), LIMITED_SUCCSS(2002), and valid message	Success
UNABLE_TO_DELIVER(3002), REALM_NOT_SERVED(3003), TOO_BUSY(3004), LOOP_DETECTED(3005), and REDIRECT_INDICATION(3006)	Failure
All other Diameter Permanent-failure Result-Code AVPs greater than and equal to 5000, and all Diameter protocol error Result-Code AVPs greater than and equal to 3000 and less than 4000	Success
Other Result-Code AVPs for invalid message or no-response	Failure

If the Result-Code value is Success, Gx notifies AAA, and AAA notifies the application that the logout is complete. If Gx does not receive a CCA-GX-T message, or if the Result-Code AVP has any other value or is missing, then the termination request is retried until the CCA-GX-T message is returned with Success. The router notifies the PCRF about subscriber logouts by sending another CCR request to be acknowledged by a CCA response. The PCRF may also use RAR requests to force subscriber logout or to change applied services.

If the Result-Code value is Failure, then the request is retried.

Subscriber Disconnect

To perform subscriber disconnects, the PCRF sends a RAR-GX-D (where D represents DISCONNECT) and the BPCEF responds with a RAA-GX-D message.

The following examples show a RAR-GX-D and RAA-GX-D packet exchange:

RAR-GX-D Packet Example

```
RAR-GX-D ::= <Diameter Header: 258, PXY, 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <string> } -- should match destination-host if configured
{ Origin-Realm:                  <string> } -- should match destination-realm
{ Destination-Realm:             <string> } -- should match origin-realm
{ Destination-Host:              <string> } -- should match origin-host
{ Re-Auth-Request-Type:          AUTHORIZE_ONLY(0) }
[ Origin-State-Id:                <u32> ] -- Indicates PCRF restart
{ Session-Release-Cause:          <enum> }
*[ Downstream data units ] -- ignored
```

RAA-GX-D Packet Example

```
RAA-GX-D ::= <Diameter Header: 272, REQ, PXY, 16777238>
{ <Session-Id> }
{ Auth-Application-Id:          16777238 }
{ Origin-Host:                  <configurable-string> }
{ Origin-Realm:                  <configurable-string> }
{ Result-Code:                   <integer> }
[ Origin-State-Id:                <u32> ]
*[ Upstream data units ]
```

The PCRF returns a RAA-GX-T message that includes the Result-Code AVP (AVP code 268) that maps to the result categories listed in [Table 88 on page 1818](#).

Table 88: Result-Code-AVP Categories

Result-Code-AVP Value	Result Category
DIAMETER_SUCCESS(2001)	Subscriber disconnect is in progress or the subscriber is not found

Table 88: Result-Code-AVP Categories *(Continued)*

Result-Code-AVP Value	Result Category
DIAMETER_UNABLE_TO_COMPLY(5012)	Subscriber is not removable
DIAMETER_TOO_BUSY(3004)	Too many outstanding disconnect requests



NOTE: The BPCEF contains buffering space for at least 512 RAR-GX-D or RAA-GX-D messages.

Connectivity Fault Recovery

Gx does not rely on the connection state between devices to detect router or PCRF outages, because some events do not affect the connection state and others are not detected when there is a Diameter relay or proxy between the devices.

To mitigate connectivity faults with the PCRF, the router uses the following fault recovery procedures:

- If the PCRF is not available, and if you installed and configured a default service, the subscriber login proceeds accordingly.
- Unacknowledged provisioning requests replay indefinitely or until the subscriber logs out.
- Logout requests wait for the final OCS interrogation to complete, and then any unacknowledged logout requests replay for 24 hours.
- The router uses standard Diameter transport redundancy to communicate with redundant PCRFs.

An important aspect of Gx fault tolerance is that subscriber login and termination requests are retried (replayed) 24 hours until a satisfactory response is received from the PCRF. You can issue the `clear network-access pcrf subscribers` command to clear all PCRF subscribers.

Understanding Gy Interactions Between the Router and the OCS

IN THIS SECTION

- [First Interrogation to the OCS | 1820](#)
- [Intermediate Interrogation to the OCS | 1823](#)

- [Final Interrogation to the OCS | 1824](#)
- [Connectivity Fault Recovery | 1826](#)
- [Abort Session Requests | 1826](#)

Information or interrogations are exchanged by means of the 3rd Generation Partnership Project (3GPP) Gy protocol between the Online Charging System (OCS), and the router acting as a Policy and Charging Enforcement Function (PCEF). Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating. PCEF optionally reports usage and receives additional authorizations from the OCS using the Gy protocol.

Starting in Junos OS Release 17.3R1, support for additional OCS and PCRF features are added using Gy and Gx protocols. The new statements:

- `accept-sdr` is added for PCRF partition at the hierarchy level `[edit access pcrf partition partition-name]`.
- `alternative-partition-name` is added for OCS partition at the hierarchy level `[edit access ocs partition partition-name]`.

Starting in Junos OS Release 18.1R1, broadband PCEF provides the file backup for OCS data when both primary and alternative paths to the OCS are not available. The CCR-GY-T frames are stored in the files on remote location. The backup is supported at the hierarchy `[edit access ocs partition partition-name]`.

After subscriber provisioning has been completed between the Policy Control and Rules Charging Function (PCRF) and PCEF, the router begins sending the following interrogations between the OCS and PCEF:

First Interrogation to the OCS

During the first interrogation, the router sends a Diameter CCR request containing a fixed set of required information to the OCS charging server. The OCS charging server then replies with validity-time, rating groups, and usage-quotas.



NOTE: For this implementation phase, the router allows subscriber access without waiting for the OCS to respond, and the OCS always grants necessary quotas.

To configure a list of charging services to communicate information with the OCS over the Gy protocol, configure the `charging-service-list ocs` statement at the `[edit access profile profile-name]` hierarchy level. The following examples show a CCR-GY-I and CCA-GY-I packet exchange:



NOTE: The T (potentially retransmitted message) bit recalculates when the CCR-GY-I is resent. This flag is set after a link failover procedure to aid the removal of duplicate requests.

CCR-GY-I Packet Example

```
CCR-GY-I ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Origin-Host:          <configurable-string> }
{ Origin-Realm:         <configurable-string> }
{ Destination-Realm:    <configurable-string> }
{ Auth-Application-Id:   4 }
{ Service-Context-Id:    98924@customer.com }
{ CC-Request-Type:       INITIAL_REQUEST(1) }
{ CC-Request-Number:     0 }
[ Destination-Host:      <configurable-string> ] -- if configured
[ User-Name:             <string> ]
[ Origin-State-Id:       <u32> ] -- if configured to send
[ Event-Timestamp:       <timestamp> ] -- login timestamp, if included by config
{ Subscription-Id:
    { Subscription-Id-Type:    <received-from-pcrf> }
    { Subscription-Id-Data:    <received-from-pcrf> }
}
{ Multiple-Services-Indicator:  MULTIPLE_SERVICES_SUPPORTED(1) }
    { Multiple-Services-CC:
        { Service-Identifier:    7 }
        { Rating-group:         292 }
    }
    { Multiple-Services-CC:
        { Service-Identifier:    7 }
        { Rating-group:         293 }
    }
    { Multiple-Services-CC:
        { Service-Identifier:    7 }
        { Rating-group:         292 }
    }
    { Multiple-Services-CC:
        { Service-Identifier:    1 }
        { Rating-group:         17 }
    }
}
```

CCA-GY-I Packet Example

```

CCA-GY-I ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Result-Code:          DIAMETER_SUCCESS(2001) }
{ Origin-Host:          <string> }          -- should match dest-host if configured
{ Origin-Realm:         <string> }          -- should match dest-realm
{ Auth-Application-Id:   4 }
{ CC-Request-Type:       INITIAL_REQUEST(1) }
{ CC-Request-Number:     0 }
{ CC-Session-Failover:   FAILOVER_NOT_SUPPORTED(0) }  -- ignored
}
{ Multiple-Services-CC:
  { Granted-Service-Unit:
    { CC-Time:           123456 }
    { CC-Total-Octets:   123455999000 }
  }
  { Service-Identifier:   7 }
  { Rating-group:         292 }
  { Validity-Time:       7200 }
  { Result-Code:          DIAMETER_SUCCESS(2001) }
}
{ Multiple-Services-CC:
  { Service-Identifier:   7 }
  { Rating-group:         293 }
  { Result-Code:          DIAMETER_CREDIT_CONTROL_NOT_APPLICABLE(4011) }
}
{ Multiple-Services-CC:
  { Service-Identifier:   7 }
  { Rating-group:         292 }
  { Result-Code:          DIAMETER_CREDIT_CONTROL_NOT_APPLICABLE(4011) }
}
{ Multiple-Services-CC:
  { Granted-Service-Unit:
    { CC-Time:           123456 }
    { CC-Total-Octets:   123455999000 }
  }
  { Service-Identifier:   1 }
  { Rating-group:         17 }
  { Result-Code:          DIAMETER_SUCCESS(2001) }
}

```

```

}
{ CC-Failure-Handling:      TERMINATE(0) }

```

Intermediate Interrogation to the OCS

After the router has sent a fixed set of required information to the OCS charging server, the OCS charging server replies with validity-time, rating groups, and usage-quotas. Validity-time and quota expirations trigger intermediate interrogation events.

Whenever a trigger event occurs on the router, an update request is sent to the OCS. The following examples show a CCR-GY-U (where U represents UPDATE_REQUEST) and CCA-GY-U packet exchange:

CCR-GY-U Packet Example

```

CCR-GY-U ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Origin-Host:      <configurable-string> }
{ Origin-Realm:     <configurable-string> }
{ Destination-Realm: <configurable-string> }
{ Auth-Application-Id: 4 }
{ Service-Context-Id: 98924@customer.com }
{ CC-Request-Type:   UPDATE_REQUEST(2) }
{ CC-Request-Number: <integer> }
[ Destination-Host:   <configurable-string> ] -- if configured
[ User-Name:          <string> ]
[ Origin-State-Id:    <u32> ] -- if configured to send
[ Event-Timestamp:    <timestamp> ] -- change timestamp, if included by config
{ Multiple-Services-Indicator:  MULTIPLE_SERVICES_SUPPORTED(1) }
{ Multiple-Services-CC:
  { Used-Service-Unit:
    { Reporting-Reason:  VALIDITY_TIME(4) }
    { CC-Time:          7200 }
    { CC-Total-Octets:  12345 }
    { CC-Input-Octets:  10000 }
    { CC-Output-Octets: 2345 }
  }
  { Service-Identifier: 7 }
  { Rating-group:       292 }
}
{ Multiple-Services-CC:
  { Used-Service-Unit:
    { Reporting-Reason:  FINAL(2) }
  }
}

```

```

    { CC-Time:          334556 }
    { CC-Total-Octets:   12345 }
    { CC-Input-Octets:   10000 }
    { CC-Output-Octets:  2345 }
  }
  { Service-Identifier:  1 }
  { Rating-group:       17 }
}
*[ More Multiple-Services-CC]

```

CCA-GY-U Packet Example

```

CCA-GY-U ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Result-Code:          DIAMETER_SUCCESS(2001) }
{ Origin-Host:          <string> } -- should match dest-host if configured
{ Origin-Realm:         <string> } -- should match dest-realm
{ Auth-Application-Id:   4 }
{ CC-Request-Type:       UPDATE_REQUEST(1) }
{ CC-Request-Number:     <integer> }
{ Multiple-Services-CC:
  { Granted-Service-Unit:
    { CC-Time:          123456 }
    { CC-Total-Octets:  123455999000 }
  }
  { Service-Identifier:   7 }
  { Rating-group:        292 }
  { Validity-Time:       7200 }
  { Result-Code:         DIAMETER_SUCCESS(2001) }
}
*[ More Multiple-Services-CC]

```

Final Interrogation to the OCS

When the client application sends a subscriber logout notice to AAA, Gy sends a CCR-GY-T (where T represents TERMINATION_REQUEST) message to notify the OCS that the provisioned subscriber is being terminated.

Whenever a trigger event occurs on the router, a terminate request is sent to the OCS. The following examples show a CCR-GY-T and CCA-GY-T packet exchange:

CCR-GY-T Packet Example

```
CCR-GY-T ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Origin-Host:          <configurable-string> }
{ Origin-Realm:         <configurable-string> }
{ Destination-Realm:    <configurable-string> }
{ Auth-Application-Id:   4 }
{ Service-Context-Id:    98924@customer.com }
{ CC-Request-Type:       TERMINATE_REQUEST(2) }
{ CC-Request-Number:     <integer> }
[ Destination-Host:      <configurable-string> ] -- if configured
[ User-Name:             <string> ]
[ Origin-State-Id:       <u32> ] -- if configured to send
[ Event-Timestamp:       <timestamp> ] -- logout timestamp, if included by config
{ Termination-Cause:     DIAMETER_LOGOUT(1) }
{ Multiple-Services-CC:
    { Used-Service-Unit:
        { Reporting-Reason:    FINAL(2) }
        { CC-Total-Octets:     12345 }
        { CC-Input-Octets:     10000 }
        { CC-Output-Octets:    2345 }
    }
    { Service-Identifier:      7 }
    { Rating-group:           292 }
}
*[ More Multiple-Services-CC]
```

CCA-GY-T Packet Example

```
CCA-GY-T ::= <Diameter Header: 272, REQ, PXY 16777238>
{ <Session-Id> }
{ Result-Code:           DIAMETER_SUCCESS(2001) }
{ Origin-Host:           <string> } -- should match dest-host if configured
{ Origin-Realm:          <string> } -- should match dest-realm
{ Auth-Application-Id:    4 }
{ CC-Request-Type:        TERMINATE_REQUEST(1) }
{ CC-Request-Number:      <integer> }
```


Connectivity Fault Recovery

Gy does not rely on the connection state between devices to detect router or OCS outages, because some events do not affect the connection state and others are not detected when there is a Diameter relay or proxy between the devices.

To mitigate connectivity faults with the OCS, the router uses the following fault recovery procedures:

- If the OCS is not available, you can configure to allow subscriber traffic by setting the `force-continue` statement at the `[edit access ocs partition partition-name]` hierarchy level.



NOTE: The `force-continue` statement is a required configuration statement.

- Unacknowledged first and intermediate interrogations replay indefinitely or until the subscriber logs out.
- Unacknowledged final interrogations replay for up to 24 hours.
- The router uses standard Diameter transport redundancy to communicate with redundant OCSs.
- You can configure transport redundancy events to trigger failures in application traffic.

An important aspect of Gy fault tolerance is that subscriber login and termination requests are retried (replayed) 24 hours until a satisfactory response is received from the OCS. You can issue the `clear network-access ocs statistics` command to clear all OCS statistics.

Abort Session Requests

The OCS may issue an ASR (Abort-Session-Request) when the receiving MX Series router collects final data and posts the final interrogation. After the MX Series router receives the response, it stops updating the OCS for the session involved.

Gy File Backup Overview

IN THIS SECTION

- [OCS SFTP-Backup | 1827](#)
- [Benefits of Gy File Backup | 1827](#)

The Gy protocol, also known as OCS, is based on incremental usage reporting while retaining the intermediate data. Therefore, the OCS server includes multiple failure protection mechanisms such as

diameter transport redundancy, alternative path to OCS, and file backup. Starting in Junos OS Release 18.1R1, broadband PCEF provides the file backup when neither primary nor alternative paths are available. The CCR-GY-T frames are stored in the files on remote location.

The OCS backup comes into effect when the OCS final-response-timeout expires. The data is queued for backup process and subscriber logout proceeds to pcrf session termination. In all cases, the backup operations are controlled by the following configuration parameters:

- **backup-limit**—limit on the total number of backup entries. After the limit is reached, new subscriber's login fails or oldest backup entries are dropped depending upon backup-overflow settings.
- **backup-timeout**— timeout for backup operation.
- **backup-overflow**—controls action when number of backup entries exceeds backup-limit.

OCS SFTP-Backup

The sftp-backup is the first backup mechanism implemented. The operations are controlled by the following parameters:

- **accumulation-timeout** – The files are written after the file accumulation time of the first CCR-GY-T submission.
- **accumulation-count** – The files are written after the number of requests are fulfilled for the file-account-count.
- **accumulation-size** – The files are written after its size is reaches the accumulation size limit.
- **retry-interval** – Every failed write operation is retried after this interval until backup-timeout is accumulated.
- **response-timeout** – The timeout on individual sftp command response.



NOTE: The OCS SFTP-Backup server is configured by its address, login, password, directory and file-prefix. A target directory exists by default, if not, the directory is created. If a target file with the same name already exists it will be overwritten.

Benefits of Gy File Backup

- Provides yet another way to deal with instability of internal network.

Understanding Interactions Between the PCRF, PCEF, and OCS


IN THIS SECTION

- [Login Interactions | 1829](#)
- [Update Interactions | 1830](#)
- [Quota Expiration and Validity-Time Interactions | 1831](#)
- [Connection and Monitoring Interactions | 1831](#)
- [Logout Interactions | 1832](#)

The Policy and Charging Rules Function (PCRF), Policy and Charging Enforcement Function (PCEF), and Online Charging System (OCS) interact to provide and charge for subscriber services. These interactions include subscriber login, service updates to existing sessions, connection and monitoring, and subscriber termination and logout.

[Figure 53 on page 1829](#) shows the components of an overall 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) architecture. The PCRF pushes the rules down to the PCEF on the MX Series router using the 3GPP Gx protocol. The PCEF provides service data flow detection, and applies the rules received from the PCRF. Optionally, the PCEF interacts with the OCS using the 3GPP Gy protocol to retrieve policy and charging authorization for quotas and credit control. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

Figure 53: 3GPP PCC Architecture Overview



The PCRF can also push changes to rules applied to an existing session. However, modifications to rating groups is not supported. You are also required to set the force-continue statement at the [edit access ocs partition *partition-name*] hierarchy level.

Login Interactions

This login sequence of events is triggered by detection of service data flow on the PCEF. This is typically a DHCP DISCOVER or PPPoE PADI packet sent by the subscriber (the CPE):

1. The PCEF sends a CCR-GX-I to the PCRF with information identifying the subscriber.
2. The PCRF replies with a CCA-GX-I to the PCEF with instructions on which rules to apply for the subscriber.
3. The PCEF installs the requested services/rules for the subscriber.
4. If OCS is being used, the PCEF sends the first interrogation to the OCS using a CCR-GY-I message, and the OCS sends applicable reports to the PCRF using a CCA-GY-I message.

If configured, the PCEF sends a notification by means of a CCR-GX-U message to the PCRF after the requested services/rules are processed.

- The rule is reported to the PCRF as *inactive* when both of the following are true:
 - The service-dynamic-profile instantiation fails.
 - Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule.

When the rule is reported as inactive, it does not affect subscriber login or other rules.

- The rule is reported to the PCRF as *active* when all of the following are true:
 - The service-dynamic-profile instantiation succeeds.
 - Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule.
 - The SUCCESSFUL_RESOURCE_ALLOCATION event trigger is set in the request.
- The report is not sent when there are no rules to report.

5. The PCRF replies back with a CCA-GX-U message.

6. The PCEF activates the services for the subscriber.

Update Interactions

This update sequence of events is triggered by an RAR-GX-U message received by the PCEF from the PCRF.

1. If the update request contains any installation or modification of rules with rating groups, then the PCEF rejects the request; otherwise, it acknowledges the request by sending an RAA-GX-U message to the PCRF.
2. The PCEF starts the service removal and installation process.
3. The PCEF waits for the service removal and installation process to complete, and if applicable, starts the final data collection process for reporting to the OCS. When the final statistics are collected, the PCEF sends a CCR-GY-U request to notify the OCS. This is a part of the removal process for an existing service in each of the following cases:
 - When the service being removed has a rating group.
 - When a new rule with a rating group was added.
 - When rules with a rating group were both removed and added.
4. The PCEF sends applicable reports to the PCRF using a CCR-GX-U message.
 - The rule is reported to the PCRF as *inactive* when both of the following are true:
 - The service-dynamic-profile instantiation fails.

- Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule.

When the rule is reported as inactive, it does not affect the update or other rules.

- The rule is reported to the PCRF as *active* when all of the following are true:
 - The service-dynamic-profile instantiation succeeds.
 - Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule.
 - The SUCCESSFUL_RESOURCE_ALLOCATION event trigger is set in the request.
- The report is not sent when there are no rules to report.

Quota Expiration and Validity-Time Interactions

For quota expirations and validity-time interactions, the router sends an intermediate interrogation to the OCS using a CCR-GY-U message and processes the OCS response.

Connection and Monitoring Interactions

When establishing a connection with the PCRF, OCS, or Diameter Relay/Proxy server, the Diameter daemon performs a standard Capability Exchange Request (CER)/Capability Exchange Answer (CEA) transaction. You use existing Junos OS Diameter infrastructure to configure an appropriate topology with the necessary redundancy features. Additionally, you can use the same Diameter connection for both PCRF and OCS communications, and other applications.

The following examples show two different communication connection scenarios:

CER Example with a Dedicated Connection Used to Communicate with the PCRF

```
CER ::= <Diameter Header: 257, REQ>
{ Origin-Realm:      CSim.PCRF.net }
{ Origin-Host:       MX-GWR3 }
{ Host-IP-Address:   10.8.52.91 }
{ Vendor-Id:         2636 }
{ Product-Name:      JUNOS }
[ Origin-State-Id:    7777 ]    -- if configured
{ Supported-Vendor-Id: 10415 }
{ Supported-Vendor-Id: 2636 }    -- have Juniper VSAs
{ Auth-Application-Id: 16777238 }
{ Vendor-Specific-Application-Id {
  { Vendor-Id:        10415 }
  { Auth-Application-Id: 16777238 }
```

```
{ Acct-Application-Id: 16777238 }
}
```



NOTE: If you set the `send-origin-state-id` statement for the router at either the `[edit access pcrf partition partition-name]` or `[edit access ocs partition partition-name]` hierarchy level, then the Origin-State-Id is included in Diameter level messages such as: CER, Device Watchdog Request (DWR)/Device Watchdog Answer (DWA), and Disconnect Peer Request (DPR)/Disconnect Peer Answer (DPA).

CER Example with a Dedicated Connection Used to Communicate with Both PCRF and OCS

```
CER ::= <Diameter Header: 257, REQ>
{ Origin-Realm:          CSim.PCRF.net }
{ Origin-Host:           MX-GWR3 }
{ Host-IP-Address:       10.8.52.91 }
{ Vendor-Id:             2636 }
{ Product-Name:          JUNOS }
[ Origin-State-Id:       7777 ] -- if configured
{ Supported-Vendor-Id:    10415 }
{ Supported-Vendor-Id:    2636 } -- have Juniper VSAs
{ Auth-Application-Id:    4 } -- this is the difference with previous
{ Auth-Application-Id:    16777238 }
{ Vendor-Specific-Application-Id {
  { Vendor-Id:            10415 }
  { Auth-Application-Id:   16777238 }
  { Acct-Application-Id:   16777238 }
}
```



NOTE: The Auth-Application-Id: 4 field and value is the authentication application ID for the OCS. This is the difference between the first and second examples.

You monitor and manage connections using standard DWR/DWA and DPR/DPA messages.

Logout Interactions

This logout sequence of events is triggered by either of the following:

- A subscriber logout request, such as a DHCP RELEASE or PPPoE PADT packet.
- The PCEF receives a RAR from the PCRF with a request to terminate a subscriber session.

The following sequence occurs when the logout is triggered:

1. The system infrastructure notifies the OCS that the subscriber logout has started.
2. If applicable, the OCS starts the final data collection process.
 - If the service being removed has a rating group, final data for this service is required to be reported. The OCS starts final data collection as necessary.
3. Both the PCRF and the PCEF wait for the final interrogation process to complete.
4. The PCEF sends a termination request (CCR-GX-T message) to the PCRF and then waits for the answer (CCA-GX-T message) from the PCRF.
5. The PCEF completes the subscriber logout process.

Understanding Upstream and Downstream Messages for the PCRF

IN THIS SECTION

- [Common Upstream Messages | 1834](#)
- [Common Downstream Messages | 1836](#)

The MX Series router implements a number of measures to protect against data overload for both downstream and upstream transactions. Downstream transactions are protected by throttling input from the network under overload conditions. The upstream transactions are protected by limiting both the number of outstanding requests and using slow retries of the first unacknowledged message for a reliable recovery.

Built-in features of the Policy and Charging Enforcement Function (PCEF) environment provide protection from overload resulting from an excessive subscriber login rate. If there are too many rule changes and disconnect operations due to Reauthorization Request (RAR-GX) messages, the router sends a Reauthorization Answer (RAA-GX) response with the result-code: DIAMETER_TOO_BUSY (3004).

Within the router's AAA component, a session represents a subscriber (client) session entry in the Session Database (SDB).



NOTE: This is a representation of subscriber session only; it is not a connection-independent permanent identifier similar to a phone number. If subscriber disconnects and reconnects, and it receives a different session ID for the second connection.

The session ID is conveyed in the Session-Id (AVP Code 263). There is a one-to-one correspondence between a session and the Session-Id value. The session ID is globally and eternally unique because it is bound to the unique router identity and used to identify a user session without any reference to other information. The same subscriber could be mapped to several sessions, such as one from a disconnect and reconnect event. However, the session is always associated with a single subscriber. The Session-Id AVP has the following default format:

```
Session-Id AVP ::= <DiameterIdentity>;
                <upper 32 bits of the AAA COMPONENT session-id>;
                <lower 32 bits of the AAA COMPONENT session-id>;
```

The *DiameterIdentity* field is the value you configure for the Diameter origin-host. Internal Session-Ids are 64-bit integers assigned in ascending order. Both numeric parts of the Session-Id string are generated using *%010u* format, which guarantees that Session-Id AVP values are in the same order lexicographically as internal 64-bit sessions.

You can also configure the router to use an extended session ID, where it appends a session stamp to the ID. The session stamp consists of the UTC time when the router creates the CCR-GX-I. In this case, the Session-Id AVP has the following format:

```
Session-Id AVP ::= <DiameterIdentity>;
                <upper 32 bits of the AAA COMPONENT session-id>;
                <lower 32 bits of the AAA COMPONENT session-id>;
                <32 bits of UTC time>;
```

The first 64 bits of the AVP remain unchanged, enabling the PCRF to trace reinitializations.

You always configure the router to use the extended session ID when it reinitializes the session in response to PCRF server errors. See ["Understanding Gx Interactions Between the Router and the PCRF" on page 1807](#) for more information.

The Policy and Charging Rules Function (PCRF) pushes rules and messages down to the PCEF using the 3GPP Gx protocol and online policy interface. The PCRF and Gx protocol include the following messages:

Common Upstream Messages

The upstream messages for Credit Control Response for Initiation, Update, and Termination (CCR-GX-I, CCR-GX-U, and CCR-GX-T) and RAA-GX may include the following rules, parameters and data:

Event-Timestamp AVP

The following shows an AVP for CCR-GX-I, CCR-GX-U, and CCR-GX-T, and RAA-GX messages between the PCRF and Gx:

```
{ Event-Timestamp: <timestamp> }
```


If you configure the Event-Timestamp AVP, it is included in the downstream message. The message definition in [Table 89 on page 1835](#) varies depending on the transaction.

Table 89: Event-Timestamp AVP Message Definition

Message	Definition
CCR-GX-I	Subscriber login timestamp

Charging Rules Installation Notifications

The following notifications show a failed installation example and a successful installation example of charging rules installation for CCR-GX-U messages between the PCRF and Gx:

**NOTE:** If unacknowledged reports are still pending at the time of the client logout, these charging rules are included in CCR-GX-T messages.

Notification Reporting a Rule Installation Failure

```
{ Charging-Rule-Report
  { Charging-Rule-Name:      <string> }
  { Charging-Rule-Name:      <string> }
  { PCC-Rule-Status:         INACTIVE(1) }
  { Rule-Failure-Code:       UNKNOWN_RULE_NAME(1) }
}
```

Notification Reporting a Rule Installation Success

```
{ Charging-Rule-Report
  { Charging-Rule-Name:      <string> }
  { Charging-Rule-Name:      <string> }
```

```
{ PCC-Rule-Status:      ACTIVE(0) }
}
```

Event Trigger Commands

The following shows a predefined event trigger command for CCR-GX and RAA-GX messages between the PCRF and Gx:

```
{ Event-Trigger: SUCCESSFUL_RESOURCE_ALLOCATION(22) }
```

Common Downstream Messages

The downstream messages for Credit Control Answer for Initiation and Update (CCA-GX-I and CCA-GX-U) and RAR-GX may include the following predefined rules with parameters and data:



NOTE: The CCA-GX-T (Termination) message is not included as a downstream message.

Charging Rule Installation Commands

The following example shows predefined rule installation commands for CCA-GX and RAR-GX messages between the PCRF and Gx:

```
{ Charging-Rule-Install
  { Charging-Rule-Name:      "fixed-cos" }
  { Charging-Rule-Definition:
    { Charging-Rule-Name:      "firewall" }
    { Service-Identifier:      10 }
    { Rating-Group:           292 }
    { Juniper-Param-V4-Input-Filter:  "my_input_filter" }
    { Juniper-Param-V4-Output-Filter:  "my_output_filter" }
  }
  [ Resource-Allocation-Notification: ENABLE_NOTIFICATION(0) ]
}
```



NOTE: Some PCRFs may be unable to generate a Resource-Allocation-Notification AVP. As a result, the `report-resource-allocation` statement at the `[edit access pcrf partition partition-name]` hierarchy level provides generated reports by default.

Charging Rule Removal Commands

The following example shows predefined rule removal commands for CCA-GX and RAR-GX messages between the PCRF and Gx:

```
{ Charging-Rule-Remove
  { Charging-Rule-Name: "predefined-ftp" }
  { Charging-Rule-Name: "firewall" }
}
```

The router processes all rule removal operations before any rule installation operations enabling you to simultaneously request both removal of an existing rule and installation of a rule having the same base name in a single transaction.

Event Trigger Commands

The following shows a predefined event trigger command for CCA-GX and RAR-GX messages between the PCRF and Gx:

```
{ Event-Trigger: SUCCESSFUL_RESOURCE_ALLOCATION(22) }
```

If the `SUCCESSFUL_RESOURCE_ALLOCATION (22)` trigger value exists in the downstream data, the Broadband PCEF reports successful installations of rules marked with Resource-Allocation-Notification AVP in the Charging-Rule-Install AVP.



NOTE: Some PCRFs may be unable to generate this event trigger. As a result, the `report-successful-resource-allocation` statement at the `[edit access pcrf partition partition-name]` hierarchy level provides generated reports by default.

Configuring the OCS Partition

The Online Charging System (OCS) works within a specific logical system: routing instance context, called a partition.



NOTE: Currently, only a single partition is supported; you must configure it within the default logical system:routing instance context.

Before you configure the OCS partition, perform the following task:

- Configure the Diameter instance at the [edit diameter] hierarchy level. See ["Configuring Diameter" on page 1762](#).

Configuration for the OCS partition consists of naming the partition and then defining or associating a numerous parameters to define the characteristics of the partition.

To configure the OCS partition:

1. Create the partition or specify the name of an existing partition.

```
[edit access ocs]
user@host# set partition partition-name
```

2. Specify the Diameter instance for the OCS partition.



NOTE: Currently, only the default Diameter instance, master, is supported.

```
[edit access ocs partition partition-name]
user@host# set diameter-instance instance-name
```

3. (Optional) Configure the Called-Station-ID AVP used in all CCR-Gy packets for the partition.

```
[edit access ocs partition partition-name]
user@host# set called-station-id station-name
```

4. (Optional) Configure the 3GPP-Charging-Id AVP used in all CCR-Gy packets for the partition.

```
[edit access ocs partition partition-name]
user@host# set charging-id number
```

5. (Optional) Configure the Destination-Host AVP value used in the CCR-GY-I message.

```
[edit access ocs partition partition-name]
user@host# set destination-host ocs-hostname
```

6. (Optional) Configure the Destination-Realm AVP value used in all CCR-GY messages

```
[edit access ocs partition partition-name]
user@host# set destination-realm ocs-realm-name
```

7. (Optional) Configure the OCS partition to the draining state to make substantial configuration changes quickly.

```
[edit access ocs partition partition-name]
user@host# set draining
```

8. (Optional) Configure the amount of time in seconds before the OCS partition responds and begins to drain after it has been placed in the draining state.

```
[edit access ocs partition partition-name]
user@host# set draining-response-timeout seconds
```

9. (Optional) Configure the amount of time in seconds before the OCS partition stops attempting to send the final interrogation during the subscriber logout process.

```
[edit access ocs partition partition-name]
user@host# set final-response-timeout seconds
```

10. Configure the OCS partition so that subscriber traffic is allowed before the first OCS interrogation and services are not removed by the PCEF when it receives negative responses from the OCS.

```
[edit access ocs partition partition-name]
user@host# set force-continue
```

11. (Optional) Configure the GGSN-Address AVP value used in all CCR-GY messages.

```
[edit access ocs partition partition-name]
user@host# set ggsn-address address
```

12. (Optional) Configure the 3GPP-GGSN-MCC-MNC AVP value used in all CCR-GY messages.

```
[edit access ocs partition partition-name]
user@host# set ggsn-mcc-mnc ggsn-mcc-mnc
```

13. (Optional) Configure the number of outstanding requests from the OCS to the OCS server that can be retried when the requests are improperly answered.

```
[edit access ocs partition partition-name]
user@host# set max-outstanding-requests number
```

14. (Optional) Specify that the Origin-State-ID AVP is included in Diameter base protocol level messages for the partition, and synchronized with the latest value sent to aid in monitoring changes in value.

```
[edit access ocs partition partition-name]
user@host# set send-origin-state-id
```

15. (Optional) Configure the information concatenated as a string in usernames that the OCS partition sends to the PCEF to identify the subscribers.
- (Optional) Include the underlying or physical interface name.

```
[edit access ocs partition partition-name]
user@host# set user-name-include base-interface-name
```

- (Optional) Use the specified character to separate the components of the username.

```
[edit access ocs partition partition-name]
user@host# set user-name-include delimiter delimiter-character
```

- (Optional) Include the specified domain name.

```
[edit access ocs partition partition-name]
user@host# set user-name-include domain-name domain-name
```

- d. (Optional) Include the interface name.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include interface-name
```

- e. (Optional) Include the client hardware MAC address from the incoming packet.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include mac-address
```

- f. (Optional) Include the NAS-Port-ID (RADIUS attribute 87) that identifies the physical interface that subscriber is using.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include nas-port-id
```

- g. (Optional) Include the name of the host that originates the Diameter message.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include origin-host
```

- h. (Optional) Include the name of the realm that originates the Diameter message.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include origin-realm
```

- i. Include the username.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include user-name
```

- j. (Optional) Include the specified prefix.

```
[edit access ocs partition partition-name]  
user@host# set user-name-include user-prefix prefix
```


16. (Optional) Configure the information required for providing file backup for OCS data.

- a. (Optional) Include the limit on the total number of backup entries for the OCS data.

```
[edit access ocs partition partition-name]  
user@host# set backup limit
```

- b. (Optional) Include the timeout for backup operation.

```
[edit access ocs partition partition-name]  
user@host# set backup timeout
```

- c. (Optional) Include the action on the number of backup entries over limit.

```
[edit access ocs partition partition-name]  
user@host# set backup overflow
```

17. (Optional) Configure the information required for providing the sftp-backup mechanism implemented for OCS data.

- a. (Optional) Configure the length of time to write the file after the first CCR-GY-T was submitted.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup accumulation-timeout
```

- b. (Optional) Configure the accumulation-count statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup accumulation-count
```

- c. (Optional) Configure the accumulation-size statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup accumulation-size
```

- d. (Optional) Configure the `retry-interval` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup retry-interval
```

- e. (Optional) Configure the `response-timeout` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup response-timeout
```

- f. (Optional) Configure the `routing-instance` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup routing-instance
```

- g. (Optional) Configure the `address` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup address
```

- h. (Optional) Configure the `port` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup port
```

- i. (Optional) Configure the `directory` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup directory
```

- j. (Optional) Configure the `file-prefix` statement to set a specific value.

```
[edit access ocs partition partition-name]  
user@host# set sftp-backup file-prefix
```

- k. (Optional) Configure the `node-ipv4-address` statement to set a specific value.

```
[edit access ocs partition partition-name]
user@host# set sftp-backup node-ipv4-address
```

- l. (Optional) Configure the `ssh-login` statement to set a specific value.

```
[edit access ocs partition partition-name]
user@host# set sftp-backup ssh-login
```

- m. (Optional) Configure the `ssh-connection-linger` statement to set a specific value.

```
[edit access ocs partition partition-name]
user@host# set sftp-backup ssh-connection-linger
```

- n. (Optional) Configure the `ssh-log-verbose` statement to set a specific value.

```
[edit access ocs partition partition-name]
user@host# set sftp-backup ssh-log-verbose
```

- o. (Optional) Configure the `ssh-passphrase` statement to set a specific value.

```
[edit access ocs partition partition-name]
user@host# set sftp-backup ssh-passphrase
```

Configuring the PCRF Partition

The Policy Control and Rules Charging Function (PCRF) works within a specific logical system: routing instance context, called a partition.



NOTE: Currently, only a single partition is supported; you must configure it within the default logical system: routing instance context.

Before you configure the PCRF partition, perform the following task:

- Configure the Diameter instance at the `[edit diameter]` hierarchy level. See ["Configuring Diameter" on page 1762](#).

Configuration for the PCRF partition consists of naming the partition and then defining or associating numerous parameters to define the characteristics of the partition.

To configure the PCRF partition:

1. Create the partition or specify the name of an existing partition.

```
[edit access pcrf]
user@host# set partition partition-name
```

2. Specify the Diameter instance for the PCRF partition.



NOTE: Currently, only the default Diameter instance, `master`, is supported.

```
[edit access pcrf partition partition-name]
user@host# set diameter-instance instance-name
```

3. (Optional) Configure the Destination-Host AVP value used in the CCR-GX-I message.

```
[edit access pcrf partition partition-name]
user@host# set destination-host pcrf-hostname
```

4. (Optional) Configure the Destination-Realm AVP value used in all CCR-GX messages

```
[edit access pcrf partition partition-name]
user@host# set destination-realm pcrf-realm-name
```

5. (Optional) Configure the PCRF to the draining state to make substantial configuration changes quickly.

```
[edit access pcrf partition partition-name]
user@host# set draining
```

6. (Optional) Configure the amount of time in seconds before the PCRF responds and begins to drain after it has been placed in the draining state.

```
[edit access pcrf partition partition-name]
user@host# set draining-response-timeout seconds
```

7. Configure the an IP connectivity access network (IP-CAN) that best fits your operating environment and access network.

```
[edit access pcrf partition partition-name]
user@host# set ip-can-type number seconds
```

8. (Optional) Configure the router to use the extended format for the session ID.



NOTE: This step is mandatory when you configure the router for local reinitialization. You might also find it useful even when you do not configure local reinitialization.



NOTE: This configuration also affects OCS sessions without any further configuration. The session ID for a given subscriber is the same for both Gx and Gy sessions.

```
[edit access pcrf partition partition-name]
user@host# set use-session-stamp
```

9. (Optional) Configure local-decision attributes for the PCRF partition to determine the behavior when the PCRF is unavailable or the PCRF does not respond in a timely manner.
 - a. (Optional) Configure subscriber login to proceed.

```
[edit access pcrf partition partition-name]
user@host# set local-decision grant
```



NOTE: You can restore the default behavior where login does not proceed by specifying deny instead of grant.

- b. (Optional) Specify how long the router waits for the PCRF to respond before using the local decision to log in the subscriber.

```
[edit access pcrf partition partition-name]
user@host# set local-decision timeout seconds
```

10. (Optional) Configure local-decision attributes for the PCRF partition so that the router reinitializes the PCRF session if the PCRF server response to the CCR-GX-I from the router is lost.



NOTE: For local reinitialization, you must also configure the following:

- The grant option
- The use-session-stamp option with the *partition* statement

- a. (Optional) Configure reinitialization to occur when the PCRF responds to a CCR-GX-I retry from the router with an unable-to-comply error code (5012) in AVP 268.

```
[edit access pcrf partition partition-name]
user@host# set local-decision reinit-on-failure
```

- b. (Optional) Configure reinitialization to occur when the PCRF erroneously responds to a CCR-GX-I retry from the router with any type of RAR.

```
[edit access pcrf partition partition-name]
user@host# set local-decision reinit-on-rar
```

- c. (Optional) Specify how long the router waits for the PCRF to respond with a CCA-GX-T before using the local decision to log in the subscriber.

```
[edit access pcrf partition partition-name]
user@host# set local-decision reinit-timeout seconds
```

11. (Optional) Configure the amount of time in seconds before the PCRF stops attempting to send a subscriber logout message.

```
[edit access pcrf partition partition-name]
user@host# set logout-response-timeout seconds
```

12. (Optional) Configure the number of outstanding requests from the PCRF to the PCRF server that can be retried when the requests are improperly answered.

```
[edit access pcrf partition partition-name]
user@host# set max-outstanding-requests number
```

13. (Optional) Specify that the PCRF sends local report downstream messages by default.

```
[edit access pcrf partition partition-name]
user@host# set report-local-rule
```

14. (Optional) Specify that the PCRF reports by default when installation fails for rules marked with the Resource-Allocation-Notification AVP in the Charging-Rule.

```
[edit access pcrf partition partition-name]
user@host# set report-resource-allocation
```

15. (Optional) Specify that the PCRF reports by default when installation either fails or succeeds for rules marked with the Resource-Allocation-Notification AVP in the Charging-Rule.

```
[edit access pcrf partition partition-name]
user@host# set report-successful-resource-allocation
```

16. (Optional) Specify that the Juniper-Dyn-Subscription-Id-Indicator AVP is included to indicate support for dynamic assignment of the subscription ID.

```
[edit access pcrf partition partition-name]
user@host# set send-dyn-subscription-indicator
```

17. (Optional) Specify that the Juniper-Network-Family-Indicator AVP is included to indicate the network families that are associated with the service request and supported by the subscriber.

```
[edit access pcrf partition partition-name]
user@host# set send-network-family-indicator
```

18. (Optional) Specify that the

```
[edit access pcrf partition partition-name]
user@host# set send-origin-state-id
```

19. (Optional) Specify that the Origin-State-ID AVP is included in Diameter base protocol level messages for the partition, and synchronized with the latest value sent to aid in monitoring changes in value.

```
[edit access pcrf partition partition-name]
user@host# set send-origin-state-id
```

20. (Optional) Configure the subscriber data to use in the PCRF partition messages to identify subscribers.

- a. (Optional) Include the underlying or physical interface name.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include base-interface-name
```

- b. (Optional) Use the specified character to separate the components of the subscription identifier.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include delimiter delimiter-character
```

- c. (Optional) Include the specified domain name.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include domain-name domain-name
```

- d. (Optional) Include the interface name.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include interface-name
```

- e. (Optional) Include the client hardware MAC address from the incoming packet.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include mac-address
```


- f. (Optional) Include the NAS-Port-ID (RADIUS attribute 87) that identifies the physical interface that subscriber is using.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include nas-port-id
```

- g. (Optional) Include the name of the host that originates the Diameter message.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include origin-host
```

- h. (Optional) Include the name of the realm that originates the Diameter message.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include origin-realm
```

- i. Include the username.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include user-name
```

- j. (Optional) Include the specified prefix.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include user-prefix prefix
```

- k. (Optional) Include the subscriber VLAN tags. You can use this instead of the `interface-name` option when the outer VLAN tag is unique across the system, which is dependent on your network topology and use case.

(Optional) Include the subscriber VLAN tags. You can use this instead of the `interface-name` option when the outer VLAN tag is unique across the system, which is dependent on your network topology and use case.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-data-include vlan-tags
```

21. (Optional) Identify the subscriber with a custom or predefined type value during the login session in CCR-GX-I and CCA-GX-I messages.

```
[edit access pcrf partition partition-name]
user@host# set subscription-id-type number
```

22. (Optional) Configure the amount of time in seconds before a PCRF partition stops attempting to send an updated rule report response using a CCR-GX-U message.

```
[edit access pcrf partition partition-name]
user@host# set update-response-timeout seconds
```

Configuring OCS Global Parameters

You can configure global attributes of the 3rd Generation Partnership Project (3GPP) Diameter credit control service charging system for the Online Charging System (OCS), which interacts with the Policy and Charging Enforcement Function (PCEF).

Currently, the only configurable global attribute is the service context identifier allocated by the service provider or operator. This value corresponds to the Service-Context-Id AVP, which together with the Service-Identifier-AVP uniquely and globally identifies the Diameter credit control service.

To configure the OCS global parameters:

- Configure the service context identifier.

```
[edit access ocs global]
user@host# set service-context-id service-context
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.2R1	Starting in Junos Release 20.1R1, you can configure the router to recover from certain PCRF server errors by reinitializing the subscriber session to resync the router and PCRF server states.

RELATED DOCUMENTATION

[Diameter Base Protocol | 1727](#)

[Gx-Plus for Provisioning Subscribers | 1781](#)

[NASREQ for Authentication and Authorization | 1852](#)

[JSRC for Subscriber Provisioning and Accounting | 1856](#)

NASREQ for Authentication and Authorization

IN THIS SECTION

- [Diameter Network Access Server Application \(NASREQ\) | 1852](#)
- [Configuring the Diameter Network Access Server Application \(NASREQ\) | 1854](#)

Diameter Network Access Server Application (NASREQ)

IN THIS SECTION

- [Benefits of Using the Diameter NASREQ Protocol | 1854](#)

The Diameter Network Access Server Requirements (NASREQ) protocol is a Diameter-based authentication, authorization, and accounting protocol defined in RFC 7155, *Diameter Network Access Server Application*. It is an alternative to using RADIUS AAA in a Diameter environment. Junos OS supports the authentication and authorization functions, but not accounting. Authentication is used for the initial subscriber login to verify the subscriber identity. Similarly, authorization is used at login to set up the initial conditions or services or both that may be needed for the subscriber. The NASREQ protocol is not used for re-authentication or re-authorization of subscribers.

Junos OS supports the following NASREQ protocol exchanges:

- AA-Request/Answer—The authentication/authorization request at login.
- Session-Termination-Request/Answer—Notification that the subscriber's session has been terminated.

- **Abort-Session-Request/Answer**—Request to terminate the subscriber's session from a NASREQ server.



NOTE: The Auth-Application-Id AVP must be set to a value of 1 in AA-Request, Session-Termination-Request, and Abort-Session-Request messages.

The NASREQ client has two queues, the transmit queue and response queue. The transmit queue stores outbound packets until they are sent to Diameter, and includes requests and responses. The response queue stores packets until Diameter responds to the request, and includes only requests waiting for a response.

The following configuration variables control transmission flow and use of the queues:

- *outstanding-requests*—The maximum number of requests (includes AAR and STR) that are sent to Diameter for wireline transmissions—effectively this is the maximum count of requests on the response-queue (the maximum number of in-flight requests for which there has not been a response or timeout); it does not include sent responses.
- *request-retry*—The number of times to re-send a given request to Diameter after it times out for its initial request. This value applies only to requests in the response queue.
- *timeout*—The number of seconds that an outbound packet remains in the transmit queue before it is declared timed out. Packets that exceed the timeout value are not transmitted. Diameter manages packets that time out after transmission. The timeout value applies to all packets in the transmit queue, including both requests and responses to be sent.

The exchange flow takes place as follows:

1. A subscriber attempts to log in and authd, acting as the NASREQ client, sends the NASREQ server a Diameter AA-Request (AAR) message that includes information about the subscriber and authentication information.
 - If the number of outstanding requests is less than the configured maximum outstanding request value, then authd sends the request to the NASREQ server for transmission and places the request on the response queue.
 - If the number of outstanding requests is greater than or equal to the configured maximum outstanding request value, then authd stores the request on the transmit queue.
2. When a response is received from the NASREQ server in the form of a Diameter AA-Answer (AAA) message, authd checks the response queue for a matching request (AAR).
 - If a matching request is found, the request is pulled from the queue and used to process the response.
 - If no matching request is found, the response is ignored and dropped.

When Diameter notifies the NASREQ client that a request has timed out, one of the following actions occurs:

- If the request is not on the response queue, the timeout is ignored.
- If the retry counter for this request is less than the configured request-retry value, authd sends the request again and increments the retry counter for that request.
- If the retry counter for this request is greater than or equal to the configured value, authd processes the request timeout and sends the next request that is on the transmit queue to the NASREQ server.

When the configured timeout period expires, authd removes any expired outbound packets from the transmit queue and processes them as having timed out.

Benefits of Using the Diameter NASREQ Protocol

- Enables the use of an external NASREQ server to provide authentication and authorization for subscribers, rather than using a RADIUS server. Some customer models might not employ a RADIUS server, or want to stop using a RADIUS server when they move to a Diameter subscriber provisioning model.

Configuring the Diameter Network Access Server Application (NASREQ)

You configure the NASREQ client as an alternative to RADIUS for subscriber authentication and authorization when the subscribers log in.

To configure NASREQ for authentication and authorization:

1. Specify NASREQ as a Diameter application (function) associated with a network element.

```
[edit diameter network-element network-element-name]
user@host# set function nasreq
```

2. Specify NASREQ as the Diameter network element forwarding function and partition.

```
[edit diameter network-element network-element-name forwarding route route]
user@host# set function nasreq
```

3. Specify NASREQ for subscriber authentication and authorization.

```
[edit access profile profile-name]
user@host# set authentication-order nasreq
```

4. Specify NASREQ for subscriber authorization only (no authentication).

```
[edit access profile profile-name]
user@host# set authorization-order nasreq
```



NOTE: When you configure both authentication-order and authorization-order, the behavior depends on the subscriber type. For DHCP subscribers, authorization-order has precedence over authentication-order. For all other subscriber types, authentication-order has precedence over authorization-order.

5. Specify the destination identity of the NASREQ partition.

```
[edit access nasreq partition partition-name]
user@host# set diameter-instance master destination-realm realm-name destination-host hostname
```

6. Specify the maximum number of requests to send to the Diameter engine for transmission. This is also the maximum number of requests in the response queue.

```
[edit access nasreq]
user@host# set max-outstanding-requests number
```

7. Specify the number of times to retry sending a request to the Diameter engine if a timeout is received from Diameter for the request.

```
[edit access [edit access nasreq]]
user@host# set request-retry retries
```

8. Specify the number of seconds an outbound packet remains in the transmit queue before it is declared timed out.

```
[edit access [edit access nasreq]]
user@host# set timeout seconds
```

RELATED DOCUMENTATION

[Diameter Base Protocol | 1727](#)

[Gx-Plus for Provisioning Subscribers | 1781](#)

[3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1799](#)

[JSRC for Subscriber Provisioning and Accounting | 1856](#)

JSRC for Subscriber Provisioning and Accounting

IN THIS SECTION

- [Juniper Networks Session and Resource Control \(SRC\) and JSRC Overview | 1856](#)
- [Understanding JSRC-SAE Interactions | 1858](#)
- [JSRC Provisioning for Dual-Stack Subscribers | 1861](#)
- [JSRC Configuration Overview | 1865](#)
- [Configuring the JSRC Partition | 1865](#)
- [Assigning a Partition to JSRC | 1866](#)
- [Authorizing Subscribers with JSRC | 1867](#)
- [Provisioning Subscribers with JSRC | 1867](#)
- [Configuring JSRC for Dual-Stack Subscribers | 1867](#)
- [Excluding AVPs from Diameter Messages for JSRC | 1868](#)
- [Service Accounting with JSRC | 1869](#)
- [Configuring Service Accounting with JSRC | 1870](#)

Juniper Networks Session and Resource Control (SRC) and JSRC Overview

IN THIS SECTION

- [Benefits of Using JSRC | 1857](#)
- [Hardware Requirements for JSRC for Subscriber Access | 1858](#)

The Juniper Networks Session and Resource Control (SRC) environment provides a central administrative point for managing subscribers and their services. The SRC software runs on Juniper

Networks C Series Controllers. The SRC software uses the Diameter protocol for communications between the local SRC peer on a Juniper Networks routing platform and the remote SRC peer on a C Series Controller. The local SRC peer is known as JSRC and is part of the AAA application. The remote SRC peer is the service activation engine (SAE); the SAE acts as the controlling agent in the SRC environment. JSRC and the SAE jointly provide the remote control enforcement functionality (RCEF).

JSRC has the following responsibilities:

- Request address authorization from the SAE.
- Request service activations from the SAE.
- Activate and deactivate services as specified by the SAE. JSRC can activate multiple policies with the same service (dynamic profile) name.
- Optionally report volume statistics for service accounting.
- Log out subscribers as specified by the SAE.
- Update the SAE with status of new service activations and deactivations.
- Synchronize subscriber state and service information with the SAE.
- Notify the SAE when subscribers log out.

The SRC software enables the SAE to activate and deactivate subscriber services (described by SRC policies) and log out subscribers. The SAE can control only those resources that have been provisioned through SAE. Therefore, the SAE receives information about only those subscribers for whom JSRC has requested provisioning from the SAE. For example, when a subscriber logs in, but the configuration did not require the session activation path to include SAE provisioning, the SAE does not receive information about this subscriber and cannot control the subscriber session.

Similarly, the SAE can control only the subscriber services that it has activated. When a service is not activated from the SAE—a RADIUS-activated service, for example—the SAE receives no information about the service and has no control over it.

The SAE can also direct JSRC to collect accounting statistics per service session.



NOTE: More than one Diameter-based application (function) can run on a router simultaneously.

Benefits of Using JSRC

- Enables the use of MX Series routers to act as the local peer in a Juniper Networks Session and Resource Control (SRC) environment for centralized management of subscribers and their services.

Hardware Requirements for JSRC for Subscriber Access

JSRC is supported on Juniper Networks MX Series 5G Universal Routing Platforms. JSRC currently supports subscriber sessions on static and dynamic interfaces.

Understanding JSRC-SAE Interactions

IN THIS SECTION

- [Subscriber Login | 1858](#)
- [Subscriber Service Activation and Deactivation | 1859](#)
- [Subscriber Resynchronization | 1859](#)
- [Subscriber Session Terminated by the SAE | 1860](#)
- [Statistics Collection and Reporting per Service Rule | 1860](#)
- [Subscriber Logout | 1860](#)

This topic describes the sequences of Diameter messages exchanged between JSRC (the local SRC peer) and the SAE (the remote SRC peer) as they interact to perform the following tasks for subscriber access:

- Subscriber login
- Service activation
- Service deactivation
- Resynchronization
- SAE-initiated subscriber logout
- Statistics collection and reporting
- Subscriber-initiated logout

Subscriber Login

JSRC authorization works as follows for different subscriber types:

- When you configure both `authorization-order jsrc` and `authentication-order` in the access profile, the authorization order applies only to DHCP subscribers affected by the profile and the authentication order applies only to non-DHCP subscribers.

- When you configure only `authorization-order jsrc` in the access profile, the authorization order applies to all subscribers affected by the profile.

When a subscriber attempts to log in, the protocol daemon sends an authentication request to AAA. In turn, JSRC sends a Diameter AA-Request message to the SAE. SAE returns a Diameter AA-Answer message that can include the Framed-IP-Address attribute and the Juniper-DHCP-Options AVP (AVP code 2010). JSRC ignores any other optional AVPs included in this AA-Answer message.

JSRC provisioning is enabled for DHCP (and SSC) subscribers when you include the `provisioning-order` statement at the `[edit access profile profile-name]` hierarchy level. When the application requests AAA to activate the subscriber's session, JSRC sends an AA-Request message that includes the Juniper-Request-Type AVP (AVP code 2050) with a value that indicates service provisioning is requested from the SAE.

The SAE returns a AA-Answer message that contains an ACK if the request is accepted or a NAK if the request is denied. If the request is accepted, the AA-Answer message includes the Juniper-Policy-Install AVP (AVP code 2020), which is used to specify the service to attach to the subscriber's interface. When this AVP is included, the SAE sets the Result-Code AVP to 1001 (DIAMETER_MULTI_ROUND_AUTH). This code means that the JSRC must send another AA-Request message to the SAE to report the success or failure of the policy instantiation (service activation) by AAA. JSRC ignores any other optional AVPs included in this AA-Answer message. The SAE returns an AA-Answer message to acknowledge this second AA-Request message.

Subscriber Service Activation and Deactivation

SAE policies provision subscriber services. After a subscriber is logged in, the SAE can send a PPR message to JSRC to activate or deactivate services. A given PPR can include the Juniper-Policy-Install AVP (AVP code 2020) to activate a service, the Juniper-Policy-Remove AVP (AVP code 2027) to deactivate a service, or both (for different services). A PPR can include no more than three of these AVPs (install, remove, or mixed).

JSRC sends a PPA message to the SAE when it has completed the tasks requested in the PPR. The PPA indicates the success or failure of the actions requested in the PPR.



NOTE: If you use RADIUS or the CLI to deactivate a service that the SAE, the SAE becomes unsynchronized with the state of subscribers on the routing engine.

Subscriber Resynchronization

During resynchronization, JSRC informs the SAE about the services that are active for the provisioned subscribers. Either JSRC or the SAE initiates the resynchronization.

- The SAE initiates resynchronization at startup or when a backup SAE takes over session control due to resource limits or conditions on the primary SAE. The SAE clears its database of all entries in preparation for the synchronization.
- JSRC initiates resynchronization at JSRC startup, such as when AAA starts or restarts.

JSRC can also initiate resynchronization in another circumstance. When an SAE in a multi-SAE environment becomes active, it must send an SRQ to JSRC as its first message. JSRC then locks the Origin-Host AVP of the active SAE. JSRC subsequently triggers resynchronization if it receives a message from any other SAE as indicated by the Origin-Host AVP. Such an incident can occur if communication between the active SAE and a standby SAE is interrupted.

Both entities initiate a resynchronization by sending an SRQ message. The recipient responds with an SRR message. After the SRR is sent, regardless of whether the SAE or JSRC initiates the synchronization, JSRC sends an AA-Request message to the SAE for each provisioned subscriber present in the session database. The AA-Request message includes a Juniper-Policy-Install AVP for the active services. The SAE returns an AA-Answer message with an ACK to acknowledge receipt.

Subscriber Session Terminated by the SAE

When the SAE terminates a subscriber session, it sends an ASR message to JSRC. JSRC causes AAA to send a logout request to the DHCP (or SSC) client application. When the DHCP client application accepts the logout request, JSRC includes an ACK in the ASR message it sends to the SAE to signify success. If the DHCP client application does not accept the request, then JSRC includes a NAK in the ASR to signify failure. The DHCP client application is responsible for initiating the actual logout sequence with AAA.

Statistics Collection and Reporting per Service Rule

Statistics information can be sent from the router to the SAE or from the SAE to the router. Both the Diameter Accounting-Request (ACR) and Accounting-Answer (ACA) messages include the Juniper-Acct-Record AVP (AVP code 2053), which identifies the policy (service) for which accounting information is requested.

Subscriber Logout

When the DHCP (or SSC) client application sends a subscriber logout notice to AAA, JSRC sends an STR message to notify the SAE that the provisioned subscriber session is being terminated. The SAE returns an STA message to JSRC, and JSRC notifies DHCP that the logout is complete.

JSRC Provisioning for Dual-Stack Subscribers

IN THIS SECTION

- [Benefits of JSRC Dual-Stack Provisioning | 1862](#)
- [AA-Request Messages When Dual-Stack Support Configured | 1862](#)
- [Accounting-Request Messages When Dual-Stack Support Configured | 1863](#)
- [Network Family Activation and Deactivation Notification When Dual-Stack Support Configured | 1864](#)

Starting in Junos OS Release 18.1R1, you can include the `dualstack-support` statement at the `[edit jsrc]` hierarchy level to configure JSRC provisioning for dual-stack subscribers so that it reports information about the separate stacks for a given subscriber, using a single JSRC session. In earlier releases, the DHCPv4 and DHCPv6 stacks are treated as a single subscriber; the remote SRC peer (SAE) is not informed about whether only one family or both families are active. The statistics are reported as an aggregate of both families rather than separated by family. The default behavior starting in Junos OS Release 18.1R1 is the same as the behavior in earlier releases.

This dual-stack provisioning behavior is not backward compatible with other releases. [Table 90 on page 1861](#) on page 2 lists some of the differences in behavior when dual-stack support is configured and when it is not configured (the default).

Table 90: Differences Between JSRC Dual-Stack Behavior by Release

Dual-Stack Support Configured	Dual-Stack Support Not Configured
When the first network family is activated, sends the addresses for only that family in the initial request to the provisioning server.	When the first network family is activated, requests provisioning from the provisioning server (SAE; remote SRC peer).
When the second network family is activated, sends a special family-activate packet to inform the provisioning server that the family is active.	When the second network family is activated, reports nothing to the provisioning server.
When the next-to-last network family is deactivated, sends a special family-deactivate packet to inform the provisioning server that the family is no longer active.	When the next-to-last network family is deactivated, reports nothing to the provisioning server.

Table 90: Differences Between JSRC Dual-Stack Behavior by Release (*Continued*)

Dual-Stack Support Configured	Dual-Stack Support Not Configured
Reports IPv4 and IPv6 statistics separately.	Reports subscriber and services statistics as an aggregate of statistics for both IPv4 and IPv6 statistics.

Benefits of JSRC Dual-Stack Provisioning

- Enables SAE to be aware of which network families are currently active for a subscriber.
- Enables collection of accurate accounting statistics per address family, rather than an aggregated count that includes statistics for both families without distinction.

AA-Request Messages When Dual-Stack Support Configured

When `dualstack-support` is configured, Diameter AA-Request (AAR) provisioning messages sent to the SAE include the following:

- IPv4 or IPv6 addresses of the currently active network families as well as the families that are in the process of being activated. When either address type is not included in the AAR message, it means that the corresponding network family is not active and is not being activated when the request is sent.
- For IPv4 addressing, the following Diameter AVPs when they are available in the subscriber's session database entry:
 - Framed-IP-Address (AVP 8)
 - Framed-IP-Netmask (AVP 9)
- For IPv6 addressing, the following Diameter AVPs and Juniper Networks Diameter AVPs when they are available in the subscriber's session database entry:
 - Framed-IPv6-Address (AVP 168)
 - Framed-IPv6-Prefix (AVP 97)
 - Delegated-IPv6-Prefix (AVP 123)
 - Juniper-IPv6-Ndra-Prefix (AVP 2200)
 - Juniper-Framed-IPv6-Netmask (AVP 2201)

- The following Juniper Networks Diameter AVPs when they are available in the subscriber's session database entry:
 - Juniper-Agent-Circuit-Id (AVP 2202)
 - Juniper-Remote-Circuit-Id (AVP 2203)
- One of the following new values in the Juniper-Request-Type AVP (2636:2050) to notify the SAE when an inactive network family activates or an existing network family deactivates:
 - 4—NETWORK_FAMILY_ACTIVATE
 - 5—NETWORK_FAMILY_DEACTIVATE

Only the addressing of the family being activated or deactivated is included in the notification.



NOTE: An activation notification is not sent for the first network family that activates. A deactivation notice is not sent for the last family that deactivates.

- When the AAR message is used for synchronization and recovery, only the addressing for the currently active address families for that subscriber. The AAR message does not include addressing for deactivated families.

Accounting-Request Messages When Dual-Stack Support Configured

When `dualstack-support` is configured, the Diameter Accounting-Request (ACR) messages always include both IPv4 and IPv6 statistics, even when the value is zero.

Statistics are reported for the life of the subscriber session and not merely for the life of the network family session. When one of the network families is inactive, JSRC continues to report the last statistics value for the inactive family with the current statistics of the active network family. If the deactivated family becomes active again, the new family statistics are added to the existing values.

The following Juniper Networks Diameter AVPs (IANA enterprise number 2636) are used to report IPv6 statistics:

- Accounting-IPv6-Input-Octets attribute (2204)
- Accounting-IPv6-Output-Octets attribute (2205)
- Accounting-IPv6-Input-Pkts attribute (2206)
- Accounting-IPv6-Output-Pkts attribute (2207)

These IPv6 AVPs are not used when `dualstack-support` is not configured. In that case the IPv6 statistics are aggregated with the IPv4 statistics in the corresponding IPv4 AVPs.

Network Family Activation and Deactivation Notification When Dual-Stack Support Configured

The following sequence describes client and authd process (daemon) behavior when a network family is activated:

1. A subscriber initiates login.
2. The client application on the router, such as PPP or DHCP, sends an authentication and authorization (AA) request to authd.
3. The authd process sends the AA request as configured and returns the response to the client application, which then returns a response to the subscriber.
4. The client application builds and configures the subscriber's interface on the router with information from the client dynamic profile.
5. The client application sends the first network family activation request to authd.
6. The authd process sends a provisioning request to the SAE that contains the addresses of the family that is being activated. Because authd sends a provisioning request for the first family activation, there is no reason to also send a family-activation notification.
7. The SAE returns policies (services) for authd to activate for the subscriber.
8. The authd process activates those services and sends a family-activation ACK response to the client application.
9. The client application might send a family activation request for the other network family.
 - a. The authd process activates any services for that specific network family and then sends a family-activation ACK response to the client application.
 - b. The authd process then sends a family-activation notification to the SAE with the addresses of the second family. The AAR message includes the Juniper-Request-Type AVP (2636:2050) with a value of 4 (NETWORK_FAMILY_ACTIVATE). The notification includes only addresses for this family.

For deactivations, the client application sends a family deactivation request only when both network families are active. The authd process deactivates the network family (and any associated services) as requested and sends the SAE a family deactivation notification with the addresses for that family. The AAR message includes the Juniper-Request-Type AVP (2636:2050) with a value of 5 (NETWORK_FAMILY_DEACTIVATE). The notification includes only addresses for this family.

However, when the last network family deactivates, the client sends a termination request, which causes authd to send the JSRC-Acct-Stop message to the SAE. Consequently there is no need for authd to send a family deactivation notification.

JSRC Configuration Overview

You can configure the JSRC client application to work with Session and Resource Control (SRC) to centrally manage subscribers and services. JSRC requests address and service authorizations from the remote SRC peer (the SAE), activates and deactivates services as specified by the SAE, logs out subscribers as specified by the SAE, and synchronizes subscriber state and service information with the SAE.

To configure JSRC:

1. Configure the JSRC partition.
See ["Configuring the JSRC Partition" on page 1865](#).
2. Assign the JSRC partition.
See ["Assigning a Partition to JSRC" on page 1866](#).
3. Configure JSRC authorization for subscribers.
See ["Authorizing Subscribers with JSRC" on page 1867](#).
4. Configure JSRC provisioning for subscribers.
See ["Provisioning Subscribers with JSRC" on page 1867](#).
5. (Optional) Configure JSRC to exclude an AVP from Messages Sent to SAE.
See ["Excluding AVPs from Diameter Messages for JSRC" on page 1868](#).
6. Configure service accounting by JSRC.
See ["Configuring Service Accounting with JSRC" on page 1870](#).
7. Configure JSRC support for dual-stack subscribers.
See ["Configuring JSRC for Dual-Stack Subscribers" on page 1867](#).
8. Configure JSRC event tracing as part of general authentication service tracing operations.
See [Tracing General Authentication Service Processes](#).

Configuring the JSRC Partition

JSRC works within a specific logical system:routing instance context, called a partition.



NOTE: Currently, only a single partition is supported; you must configure it within the default logical system:routing instance context.

Before you configure the JSRC partition, perform the following task:

- Configure the Diameter instance at the [edit diameter] hierarchy level. See ["Configuring Diameter" on page 1762](#).

Configuration for the JSRC partition consists of naming the partition and then associating a Diameter instance, the SAE hostname, and the SAE realm with the partition.

To configure the JSRC partition:

1. Create the partition.

```
[edit jsrc]
user@host# set partition partition1
```

2. Specify the Diameter instance for the JSRC partition.



NOTE: Currently, only the default Diameter instance, `master`, is supported.

```
[edit jsrc partition partition1]
user@host# set diameter-instance master
```

3. Configure the destination host for the JSRC partition.

```
[edit jsrc partition partition1]
user@host# set destination-host sae1
```

4. Configure the destination realm for the JSRC partition.

```
[edit jsrc partition partition1]
user@host# set destination-realm generic.example.com
```

Assigning a Partition to JSRC

You must associate a configured JSRC partition with the JSRC instance that you are configuring.

Before you assign a partition to JSRC, perform the following task:

- Configure the JSRC partition. See ["Configuring the JSRC Partition" on page 1865](#)

To assign the JSRC partition:

- Specify the partition name.

```
[edit jsrc]
user@host# set jsrc-partition partition1
```

Authorizing Subscribers with JSRC

You can configure AAA to use JSRC in an SRC environment to request authorization from the SAE when AAA is verifying whether a DHCP subscriber can access the router. When JSRC authorization is configured, AAA ignores any configured authentication order settings.

Before you configure JSRC authorization, perform the following tasks:

- Create the subscriber access profile at the [edit access profile] hierarchy level.
- Define the subscriber username with the `username-include` statement in the authentication configuration for DHCP local server or DHCP relay.

To configure JSRC authorization:

- Specify `jsrc` as the authorization method in the profile.

```
[edit access profile dhcpsub1]
user@host# set authorization-order jsrc
```

Provisioning Subscribers with JSRC

You can configure AAA to use JSRC in an SRC environment to request provisioning from the SAE to instantiate services for an authenticated subscriber.

Before you configure JSRC provisioning for subscribers, perform the following task:

- Create the subscriber access profile at the [edit access profile] hierarchy level.

To configure JSRC provisioning:

- Specify `jsrc` as the provisioning method in the profile.

```
[edit access profile dhcpsub1]
user@host# set provisioning-order jsrc
```

Configuring JSRC for Dual-Stack Subscribers

By default, JSRC provisioning for dual-stack subscribers treats the DHCPv4 and DHCPv6 stacks as a single subscriber. The remote SRC peer (SAE) is not informed about whether only one family or both families are active. The statistics are reported as an aggregate of both families rather than separated by family.

Starting in Junos OS Release 18.1R1, you can configure dual-stack support so that JSRC reports information about the separate stacks for a given subscriber, using a single JSRC session.

When you configure dual-stack support for JSRC, Diameter AA-Request (AAR) provisioning messages sent to the SAE include Diameter AVPs (IANA enterprise number 2636) to convey the IPv4 and IPv6 addressing information that is available in the session database.

For IPv4, that includes the following AVPs:

- Framed-IP-Address (AVP 8)
- Framed-IP-Netmask (AVP 9)

For IPv6, that includes the following AVPs:

- Framed-IPv6-Address (AVP 168)
- Framed-IPv6-Prefix (AVP 97)
- Delegated-IPv6-Prefix (AVP 123)
- Juniper-IPv6-Ndra-Prefix (AVP 2200)
- Juniper-Framed-IPv6-Netmask (AVP 2201)

JSRC also includes information about the access line if it is available in the session database, by means of Juniper-Agent-Circuit-Id (AVP 2202) and Juniper-Remote-Circuit-Id (AVP 2203).

When the first network family is activated, JSRC sends the addresses for only that family in the initial request to the provisioning server. When the second network family is activated, the AAR message includes the Juniper-Request-Type AVP (2050) with a value of 4 to signify family activation. When the next-to-last family is deactivated, the same AVP is sent with a value of 5 to signify the deactivation.

To configure JSRC provisioning to report dual-stack subscriber information by family:

- Enable dual-stack support.

```
[edit jsrc]
user@host# set dualstack-support
```

Excluding AVPs from Diameter Messages for JSRC

Starting in Junos OS Release 14.2, you can configure the router to exclude AVPs from Diameter messages that are sent to the SAE from JSRC.



NOTE: Currently, only the user-name (1) AVP is supported.

To configure JSRC to exclude an AVP in Diameter messages:

1. Specify that you want to configure JSRC settings in the access profile.

```
[edit access profile profilewestern55]
user@host# edit jsrc
```

2. Specify that you want to configure Diameter attribute usage.

```
[edit access profile profilewestern55 jsrc]
user@host# edit attributes
```

3. Configure the router to exclude the specified AVP from the specified messages. The following example excludes the user-name AVP from authorization and provisioning AAR messages.

```
[edit access profile profilewestern55 jsrc attributes]
user@host# set exclude user-name authorization-request
user@host# set exclude user-name provisioning-request
```

Service Accounting with JSRC

A service session represents a service for a specific subscriber. Service sessions exist in the context of a subscriber session. JSRC activates and deactivates services as specified by the SAE (remote SRC peer). JSRC can collect and report service accounting data by volume. JSRC accounting requires that either classic firewall filters or fast update firewall filters be configured to count service packets—the service packet information provides the volume statistics.



NOTE: JSRC supports only volume statistics accounting for service sessions. Time statistics and subscriber accounting are not supported.

JSRC service accounting supports both accounting based on service activation/deactivation and interim accounting.

- Service activation/deactivation accounting—When accounting is enabled, JSRC sends an accounting start message to the SAE when it activates a service and an accounting stop message when it deactivates the service. The start message initiates the accounting session and provides initial

information about the service session. The stop message terminates the accounting session and reports the final (cumulative) accounting data.

- Interim accounting—When interim accounting is enabled for a service session, JSRC sends interim accounting messages to the SAE at a specified interval to report the cumulative accounting information available at that time. Interim accounting is ignored when accounting is not enabled for the corresponding service session.

JSRC accounting for a service begins when the service is activated, and remains in effect while the service is active. The SAE specifies the service (policy) to be activated for the subscriber with the Juniper-Policy-Install AVP (AVP code 2020). When this AVP includes the Juniper-Acct-Collect AVP (AVP code 2054), JSRC initiates service activation/deactivation accounting for the service.

JSRC initiates interim accounting when the Juniper-Policy-Install AVP includes the Acct-Interim-Interval AVP (AVP code 85). In this case, JSRC updates the accounting values at the interval specified in the AVP — in the range 600 through 86,400 seconds. Aggregate counters are reported for the dual stack case.

JSRC and the SAE exchange Diameter Accounting-Request (ACR) and Accounting-Answer (ACA) messages to communicate accounting data. Both messages include the Juniper-Acct-Record AVP (AVP code 2053) to identify the service for which accounting information is requested.

JSRC sends ACR messages to report accounting data to the SAE. The ACR message includes the Accounting-Record-Type AVP (AVP code 480) to specify the kind of accounting record that it is sending. When a service is activated, this AVP has a value of START_RECORD. When a service is deactivated, it has a value of STOP_RECORD. For interim accounting, ACR messages are sent at the specified accounting interval and the AVP has a value of INTERIM_RECORD.

In addition to specifying the accounting record type, the ACR messages include standard RADIUS attributes to specify the desired statistics: Acct-Input-Octets [42], Acct-Output-Octets [43], Acct-Input-Packets [47], Acct-Output-Packets [48], and Acct-Session-Time [46].

The SAE returns ACA messages to the JSRC to acknowledge receipt of the ACR messages.

An access profile specifies subscriber access authentication and accounting parameters. When a service is activated through JSRC, the accounting reports can be sent either to the SAE or to RADIUS. The default configuration sends the reports to the SAE; you can also configure this by including the `service accounting-order activation-protocol` statement in the access profile. To send the reports instead to the RADIUS server, include the `service accounting-order radius` statement in the access profile.

When a service is activated through RADIUS rather than through JSRC, the accounting reports of the service session are sent to the RADIUS server.

Configuring Service Accounting with JSRC

In addition to the configuration shown here, the network context for JSRC service accounting includes the configuration of firewall filters to count the statistics, Diameter, JSRC, the subscriber services, RADIUS, and the SRC.

You can configure JSRC to report accounting statistics for service sessions.

To configure service accounting by JSRC:

1. Configure JSRC to provision subscriber services.

```
[edit access profile profile-name
user@host# set provisioning-order jsrc
```

2. Configure service accounting to be provided by the application that provisions the service—JSRC.

```
[edit access profile profile-name service]
user@host# set accounting-order activation-protocol
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, you can include the dualstack-support statement at the [edit jsrc] hierarchy level to configure JSRC provisioning for dual-stack subscribers so that it reports information about the separate stacks for a given subscriber, using a single JSRC session.
14.2	Starting in Junos OS Release 14.2, you can configure the router to exclude AVPs from Diameter messages that are sent to the SAE from JSRC.

RELATED DOCUMENTATION

Diameter Base Protocol 1727
Gx-Plus for Provisioning Subscribers 1781
3GPP Policy and Charging Control for Wireline Provisioning and Accounting 1799
NASREQ for Authentication and Authorization 1852
JSRC and Subscribers on Static Interfaces 1872

JSRC and Subscribers on Static Interfaces

IN THIS SECTION

- [Subscribers on Static Interfaces Overview | 1872](#)
- [Subscribers over Static Interfaces Configuration Overview | 1876](#)
- [Example: Configuring Static Subscribers for Subscriber Access | 1877](#)
- [Specifying the Static Subscriber Global Access Profile | 1879](#)
- [Specifying the Static Subscriber Global Dynamic Profile | 1879](#)
- [Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers | 1880](#)
- [Configuring the Static Subscriber Global Authentication Password | 1880](#)
- [Configuring the Static Subscriber Global Username | 1881](#)
- [Creating a Static Subscriber Group | 1883](#)
- [Specifying the Static Subscriber Group Access Profile | 1884](#)
- [Specifying the Static Subscriber Group Dynamic Profile | 1884](#)
- [Specifying the Static Subscriber Group Service Profile | 1884](#)
- [Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group | 1885](#)
- [Configuring the Static Subscriber Group Authentication Password | 1885](#)
- [Configuring the Static Subscriber Group Username | 1886](#)

Subscribers on Static Interfaces Overview

IN THIS SECTION

- [Benefits of Subscribers on Static Interfaces | 1875](#)

You can associate subscribers with statically configured interfaces and provide dynamic service activation and deactivation for these subscribers. When the static interface comes up, the event is treated as a subscriber login. When the interface goes down, it is treated as a subscriber logout.

You can configure the static subscribers to be authenticated and authorized by means of RADIUS. In this case, RADIUS can then activate and deactivate services with change of authorization (CoA) messages.

However, this configuration does not prevent the interface from coming up and forwarding traffic. Further, authorization parameters are not imposed on the subscriber interface.

Alternatively, you can use JSRC for dynamic service activation and deactivation for these subscribers. After the subscribers are present in the session database (SDB), JSRC can report the subscribers to the SAE so that the SRC software can subsequently manage the subscribers.

The following guidelines apply to static subscribers:

- Static subscribers are supported only on Ethernet interfaces, static demux interfaces, and pseudowire interfaces over logical tunnels (PS/LT). PS/LT support, introduced in Junos OS Release 18.3R1, enables full subscriber management (equivalent to dynamic subscribers) for statically provisioned subscribers whose traffic is transported over IP/MPLS access models.
- Only one static subscriber can exist over a given interface.
- An interface cannot appear in more than one group.
- Static subscribers cannot be created over dynamic interfaces.

Static subscribers are intended to work with JSRC. Include the provisioning-order `jsrc` statement at the `[edit access profile profile-name]` hierarchy level to enable JSRC to handle the subscribers at the direction of the SRC software.

If the authentication request fails for a static subscriber, a 60-minute, nonconfigurable timer begins counting down. The request is reissued when the timer expires. This action repeats for as long as the interface is operationally up.

You can force a logout of the static subscriber by issuing the request `services static-subscribers logout interface interface-name` command. A static subscriber can also be logged out by AAA or an external policy manager. In both cases, no subsequent logins can take place on the underlying interface until you reset the state by issuing the request `services static-subscribers login interface interface-name` command or the router or process reboots.

You can log out an interface group by issuing the request `services static-subscriber logout group group-name` command. You can subsequently log in a group of interfaces by issuing the request `services static-subscriber login group group-name` command.

No new CLI statements are required to configure the dynamic profile for static subscribers. The dynamic profile can be very simple; it is activated at login and deactivated at logout. If you do not configure a profile, then the *junos-default-profile* is automatically activated.

During a *graceful Routing Engine switchover* (GRES) event, active static subscribers are recovered, inactive subscribers are cleaned up, and logout continues for subscribers that were in the process of logging out.

Include the `static-subscribers` statement at the `[edit system services]` hierarchy level to configure static subscribers. Include the `traceoptions` statement at the `[edit system processes static-subscribers]` hierarchy level to configure tracing operations for static subscribers.

You can configure the access profile, dynamic profile, service profile, and authentication parameters for all static subscribers or for a particular group of static subscribers:

- To configure the access profile that triggers AAA services for the static subscriber for all static subscribers, include the `access-profile` statement at the `[edit system services static-subscribers]` hierarchy level. Alternatively, include this statement at the `[edit system services static-subscribers group group-name]` hierarchy level to apply the profile to a specific group and override a top-level configuration.
- To configure the dynamic profile that is instantiated when the static subscriber logs in for all static subscribers, include the `dynamic-profile` statement at the `[edit system services static-subscribers]` hierarchy level. Alternatively, include this statement at the `[edit system services static-subscribers group group-name]` hierarchy level to apply the profile to a specific group and override a top-level configuration. Do not specify a dynamic profile that creates a dynamic interface.
- To configure the service profile for all static subscribers at the global level and at the group level, include the `service-profile` statement at the `[edit system services static-subscribers group group-name]` hierarchy level.
- To configure the authentication parameters that trigger an Access-Request message to AAA for all static subscribers, include the `authentication` statement at the `[edit system services static-subscribers]` hierarchy level. Alternatively, include the statement at the `[edit system services static-subscribers group group-name]` hierarchy level to configure authentication for a specific group and override a top-level configuration. If you do not configure authentication, then by default the interface name is modified and used as the default username for the subscriber session and the authentication request.

The configurable authentication parameters include the password and details of how the username is formed. Include the `password` statement at the `[edit system services static-subscribers authentication]` hierarchy level to configure the authentication password for all static subscribers. Alternatively, include the statement at the `[edit system services static-subscribers group group-name authentication]` hierarchy level to configure authentication for a specific group and override a top-level configuration.

The username that is sent to AAA for authentication must include at least one of the following attributes:

- Domain name
- User prefix
- Interface name
- Logical system name

- Routing instance name

To configure how the username is formed for all static subscribers, include the desired statements at the [edit system services static-subscribers authentication] hierarchy level: domain-name, user-prefix, logical-system-name, or routing-instance-name. Alternatively, include the desired statements at the [edit system services static-subscribers group *group-name* authentication] hierarchy level to configure the username for a specific group and override a top-level configuration.

If you change the authentication configuration for an existing group or for static subscribers globally, the change has no effect on existing static subscribers. The changes are applied only to any new logins that are attempted after you commit the changes.

A group configuration must specify all the interfaces that you expect to support static subscribers. Include the interface statement at the [edit system services static-subscribers group *group-name*] hierarchy level to specify the interfaces. This statement enables you to specify a single interface or a range of interfaces.

You must also statically configure these interfaces before any static subscribers can be supported on them. You must configure the static interfaces in the same logical system and routing instance as the group that includes the interfaces.

If you change the interfaces that are included in an existing interface group, existing static subscribers are automatically logged out and then back in when you commit the changes. However, changes made to the configuration of the interface itself have no effect on the login or logout state of the static subscriber associated with that interface.

By default, multiple subscribers are not supported on top of the same VLAN *logical interface*. If you want to support this behavior, then you can manage multiple subscribers on a single logical interface in one of two ways. You can either merge attributes such as firewall filters and CoS attributes for the multiple subscribers, or you can replace the current attributes with those of a new subscriber whenever a new subscriber logs into the underlying VLAN logical interface.

- To enable attribute merging for all static interfaces, include the aggregate-clients merge statement at the [edit system services static-subscribers] hierarchy level. Alternatively, include this statement at the [edit system services static-subscribers group *group-name*] hierarchy level to enable attribute merging for a specific group of static interfaces and override a top-level configuration.
- To enable attribute replacement for all static interfaces, include the aggregate-clients replace statement at the [edit system services static-subscribers] hierarchy level. Alternatively, include this statement at the [edit system services static-subscribers group *group-name*] hierarchy level to enable attribute replacement for a specific group of static interfaces and override a top-level configuration.

Benefits of Subscribers on Static Interfaces

- Offers static-subscribers the ability to configure service-profile.

- Provides dynamic service activation for the associated subscribers with statically configured interfaces.
- Provides competitive advantage with RFC compliancy.

Subscribers over Static Interfaces Configuration Overview

This topic describes the procedure for configuring subscribers over static interfaces (static subscribers).

Before you configure subscribers over static interfaces, perform the following tasks:

- Configure the static interfaces on which you want to create and manage subscribers.
- Create an access profile to trigger AAA services for static subscribers.
- Create a dynamic profile that is instantiated when static subscribers log in.

To configure static subscribers:

1. Specify the global access profile that triggers AAA services for static subscribers.
See ["Specifying the Static Subscriber Global Access Profile" on page 1879](#).
2. Specify the global dynamic profile that is instantiated when static subscribers log in.
See ["Specifying the Static Subscriber Global Dynamic Profile" on page 1879](#).
3. Configure global method to handle multiple subscribers on a VLAN Logical Interface.
See ["Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers" on page 1880](#).
4. Configure the global authentication password for static subscribers.
See ["Configuring the Static Subscriber Global Authentication Password" on page 1880](#).
5. Configure the global username for static subscribers.
See ["Configuring the Static Subscriber Global Username" on page 1881](#).
6. Configure a group of subscribers to share values different from the global configuration.
See ["Creating a Static Subscriber Group" on page 1883](#).
7. Specify the access profile for the static subscriber group.
See ["Specifying the Static Subscriber Group Access Profile" on page 1884](#).
8. Specify the dynamic profile for the static subscriber group.
See ["Specifying the Static Subscriber Group Dynamic Profile" on page 1884](#).
9. Specify the service profile for the static subscriber group.
See ["Specifying the Static Subscriber Group Service Profile" on page 1884](#).
10. Configure method to handle multiple subscribers on a VLAN Logical Interface for a static subscriber group.

See ["Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group"](#) on page 1885.

11. Configure the authentication password for the static subscriber group.
See ["Configuring the Static Subscriber Group Authentication Password"](#) on page 1885.
12. Configure the username for the static subscriber group.
See ["Configuring the Static Subscriber Group Username"](#) on page 1886.
13. (Optional) Force a static subscriber to be logged out from an interface.
See ["Forcing a Static Subscriber to Be Logged Out"](#) on page 1901.
14. (Optional) Enable an interface to accept static subscriber logins.
See ["Resetting the State of an Interface for Static Subscriber Login"](#) on page 1901.
15. (Optional) Force static subscribers to be logged out from a group of interfaces.
See ["Forcing a Group of Static Subscribers to Be Logged Out"](#) on page 1902.
16. (Optional) Enable a group of interfaces to accept static subscriber logins.
See ["Resetting the State of an Interface Group for Static Subscriber Login"](#) on page 1902.
17. Configure trace options for troubleshooting the configuration.
See ["Tracing Static Subscriber Events for Troubleshooting"](#) on page 1903.

Example: Configuring Static Subscribers for Subscriber Access

This example shows a static subscriber configuration.

1. Configure the access profile to be used for static subscribers.

```
access {
  profile access5 {
    provisioning-order jsrsc;
    accounting {
      order radius;
    }
    authentication {
      order radius;
    }
  }
}
```

2. Configure the dynamic profile to be used for static subscribers.

If you do not configure this profile, the default profile, junos-default-profile, is used.

3. Configure the static interfaces on which to layer the static subscribers.
4. Configure the parameters that apply globally to all static subscribers in the configuration context.

```
static-subscribers {
    access-profile access5;
    dynamic-profile dyn-profile-1;
    authentication {
        password $ABC123;
        username-include {
            user-prefix Building5;
            interface;
            logical-system-name;
            routing-instance-name;
            domain-name example.com;
        }
    }
}
```

5. If you want to override the global parameters for certain static subscribers, create a group of static interfaces for those subscribers and configure parameters to apply to that group. Repeat this step for as many groups as you need.

```
static-subscribers {
    group boston {
        interface ge-1/0/1.1 upto ge-1/0/1.102
        interface ge-1/0/1.6 exclude
        interface ge-1/0/1.70 upto ge-1/0/1.80 exclude
        access-profile boston-acs;
        dynamic-profile dyn-profile-2;
        authentication {
            password $ABC123;
            username-include {
                user-prefix 2ndFloor;
                interface;
                logical-system-name;
                routing-instance-name;
                domain-name example.net;
            }
        }
    }
}
```

```
    }
}
```

6. Configure tracing options for static subscriber events.

```
static-subscribers {
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Specifying the Static Subscriber Global Access Profile

You specify a previously created access profile that triggers AAA services for all static subscribers. This value can be overridden for a group of static subscribers when a different profile is configured for that group.

To specify the access profile used for all static subscribers:

- Specify the profile name.

```
[edit system services static-subscribers]
user@host# set access-profile access5
```

Specifying the Static Subscriber Global Dynamic Profile

You specify a previously created dynamic profile that is instantiated when a static subscriber logs in. This profile is used for all static subscribers. This value can be overridden for a group of static subscribers when a different profile is configured for that group.

To specify the dynamic profile used for all static subscribers:

- Specify the profile name.

```
[edit system services static-subscribers]
user@host# set dynamic-profile dyn-profile-1
```

Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers

For a given interface, only a single static subscriber (or group) is logged in. Although we do not recommend this practice, you might have other kinds of subscribers configured on the same interface, such as a DHCP subscriber managed by the DHCP application. You can use the `aggregate-clients` statement to extend the dynamic profile for all static subscribers to enable multiple subscribers to share the same VLAN logical interface.

You can specify that attributes (such as CoS or firewall) for the multiple subscribers are merged for the logical interface. That is, the profiles for multiple subscribers of different types are instantiated on the interface, but the profile attributes of each are merged together. Alternatively, you can specify that the instantiated profile for the current subscriber is replaced by the profile of a new subscriber that logs in using the same logical interface. This configuration can be overridden for a group of static subscribers when a different configuration is applied for that group.



NOTE: The `aggregate-clients` statement is not supported for enhanced subscriber management.

To enable multiple subscribers to share the same VLAN logical interface for all static subscribers, do one of the following:

- Specify that the multiple subscriber attributes are merged for the logical interface.

```
[edit system services static-subscribers dynamic-profile dyn-profile-1]
user@host# set aggregate-clients merge
```

- Specify that the entire logical interface is replaced when a new subscriber logs into the network using the same VLAN logical interface.

```
[edit system services static-subscribers dynamic-profile dyn-profile-3]
user@host# set aggregate-clients replace
```

Configuring the Static Subscriber Global Authentication Password

You configure a password that is included in the Access-Request message sent to AAA to authenticate all static subscribers. This value can be overridden for a group of static subscribers when a different password is configured for that group.

To specify the authentication password used for all static subscribers:

- Specify the password.

```
[edit system services static-subscribers authentication]
user@host# set password $ABC123
```

Configuring the Static Subscriber Global Username

You configure how the username is formed. The username serves as the username for all static subscribers that are created and is included in the Access-Request message sent to AAA to authenticate all static subscribers. This value can be overridden for a group of static subscribers when a different username is configured for that group.

The username must include at least one of the possible elements. The value of each element is concatenated in a specific order; the resulting string is the username. If you specify their inclusion, the interface name, logical system name, routing instance name, and VLAN tags are derived from the configuration context. The elements are ordered as follows (shown with the default delimiter):

*user-prefix.interface.outer-tag-inner-tag.logical-system-name.
routing-instance-name@domain-name*

To configure the username for all static subscribers:

1. (Optional) Specify a prefix for the username.

```
[edit system services static-subscribers authentication username-include]
user@host# set user-prefix user-prefix-string
```

2. (Optional) Specify that the interface name is included in the username.

```
[edit system services static-subscribers authentication username-include]
user@host# set interface
```

3. (Optional) Specify that the VLAN tags (VLAN IDs) associated with the static interface are included in the username. For single-tagged VLANs, the component is the *outer-tag*. For dual-tagged (stacked) VLANs, the component is *outer-tag-inner-tag*. For IP demux interfaces configured for static subscribers, the VLAN tags configured on the underlying interface are used.

```
[edit system services static-subscribers authentication username-include]
user@host# set vlan-tags
```


4. (Optional) Specify that the logical system name is included in the username.

```
[edit system services static-subscribers authentication username-include]
user@host# set logical-system-name
```

5. (Optional) Specify that the routing instance name is included in the username.

```
[edit system services static-subscribers authentication username-include]
user@host# set routing-instance-name
```

6. (Optional) Specify a domain name to include in the username.

```
[edit system services static-subscribers authentication username-include]
user@host# set domain-name domain-name
```

7. (Optional) Specify a delimiter character to separate the username elements except for the domain name. The domain name is always preceded by the @ character. The default delimiter is a period (.)

```
[edit system services static-subscribers authentication username-include]
user@host# set delimiter delimiter-character
```

Consider the following configuration:

```
[edit system services static-subscribers authentication username-include]
user@host# set user-prefix Building5
user@host# set interface
user@host# set logical-system-name
user@host# set routing-instance-name
user@host# set domain-name campus.example.com
```

Configured in the default logical system and master routing instance for interface ge-0/1/1.100, this sample configuration generates the following username:

Building5.ge-0-1-1-100.default.master@campus.example.com

Now consider a different configuration, where the static interface has a dual-tagged VLAN, with an outer VLAN ID of 4040 and an inner VLAN ID of 3000:

```
[edit system services static-subscribers authentication username-include]
user@host# set user-prefix Floor12
```

```

user@host# set domain-name Bldg5.example.com
user@host# set vlan-tags
user@host# set delimiter $

```

This sample configuration generates the following username:

```
Floor12$4040-3000@Bldg5.example.com
```

Even though a delimiter of \$ is configured, outer and inner VLAN IDs are always separated by - and the domain name is always separated from preceding elements by @.

Creating a Static Subscriber Group

You can override the configuration that is applied globally to static subscribers by creating a static subscriber group that consists of a set of statically configured interfaces. You can then apply a common configuration for the group with values different from the global values for access and dynamic profiles, password, and username.

To configure an interface group for static subscribers:

1. Access the [edit system services static-subscribers] hierarchy level.
2. Create the group and assign the name.

```

[edit system services static-subscribers]
user@host# edit group boston

```

3. Specify the names of one or more interfaces on which static subscribers can be created. You can repeat the *interface interface-name* statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```

[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.1
user@host# set interface ge-1/0/1.2

```

4. (Optional) You can use the *upto upto-interface-name* option to specify a range of interfaces for a group.

```

[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.3 upto ge-1/0/1.9

```

5. (Optional) You can use the `exclude` option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.1 upto ge-1/0/1.102
user@host# set interface ge-1/0/1.6 exclude
user@host# set interface ge-1/0/1.70 upto ge-1/0/1.80 exclude
```

Specifying the Static Subscriber Group Access Profile

You can override the configured global access profile by specifying a different profile for a group of static subscribers. The access profile triggers AAA services for that group of static subscribers.

To specify the access profile used for a group of static subscribers:

- Specify the profile name.

```
[edit system services static-subscribers group boston]
user@host# set access-profile boston-acs
```

Specifying the Static Subscriber Group Dynamic Profile

You can override the configured global dynamic profile by specifying a different profile for a group of static subscribers. The dynamic profile is instantiated when any static subscriber in the group logs in.

To specify the dynamic profile used for a group of static subscribers:

- Specify the profile name.

```
[edit system services static-subscribers group boston]
user@host# set dynamic-profile dyn-profile-2
```

Specifying the Static Subscriber Group Service Profile

When external policy server is unavailable, you can assign a default dynamic service profile to be applied to a static subscriber session by specifying the service profile from Junos OS Release 17.4R1 onwards. The service profile can be specified at the group level and at the global level. Specify `service-profile` statement at the `[edit system services static-subscribers group group-name]` hierarchy level

To specify the service profile used for a group of static subscribers:

- Specify the dynamic service profile name.

```
[edit system services static-subscribers group group-name]
user@host# set service-profile service-profile-name
```

Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group

For a given interface, only a single static subscriber group (or static subscriber) is logged in. Although we do not recommend this practice, you might have other kinds of subscribers configured on the same interface, such as a DHCP subscriber managed by the DHCP application. You can use the `aggregate-clients` statement to extend the dynamic profile for a static subscriber group to enable multiple subscribers to share the same VLAN logical interface.

You can specify that attributes (such as CoS or firewall) for the multiple subscribers are merged for the logical interface. That is, the profiles for multiple subscribers of different types are instantiated on the interface, but the profile attributes of each are merged together. Alternatively, you can specify that the instantiated profile for the current subscriber group is replaced by the profile of a new subscriber that logs in using the same logical interface. This configuration overrides the configuration applied to all static subscribers that are not members of the group.

To enable multiple subscribers to share the same VLAN logical interface for a static subscriber group, do one of the following:

- Specify that the multiple subscriber attributes are merged for the logical interface.

```
[edit system services static-subscribers group boston dynamic-profile dyn-profile-2]
user@host# set aggregate-clients merge
```

- Specify that the entire logical interface is replaced when a new subscriber logs into the network using the same VLAN logical interface.

```
[edit system services static-subscribers group boston dynamic-profile dyn-profile-4]
user@host# set aggregate-clients replace
```

Configuring the Static Subscriber Group Authentication Password

You can override the configured global authentication password by specifying a different password for a group of static subscribers. This password is included in the Access-Request message sent to AAA to authenticate all static subscribers in the group.

To specify the authentication password used for a group of static subscribers:

- Specify the password.

```
[edit system services static-subscribers group boston authentication]
user@host# set password $ABC123
```

Configuring the Static Subscriber Group Username

You can override the configured global username by specifying a different username for a group of static subscribers. The username serves as the username for a group of static subscribers that is created and is included in the Access-Request message sent to AAA to authenticate that group.

The username must include at least one of the possible elements. The value of each element is concatenated in a specific order; the resulting string is the username. If you specify their inclusion, the interface name, logical system name, routing instance name, and VLAN tags are derived from the configuration context. The elements are ordered as follows (shown with the default delimiter):

*user-prefix.interface.outer-tag-inner-tag.logical-system-name.
routing-instance-name@domain-name*

To configure the username for a group of static subscribers:

1. (Optional) Specify a prefix for the username.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set user-prefix user-prefix-string
```

2. (Optional) Specify that the interface name is included in the username.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set interface
```

3. (Optional) Specify that the VLAN tags (VLAN IDs) associated with the static interface are included in the username. For single-tagged VLANs, the component is the *outer-tag*. For dual-tag (stacked) VLANs, the component is the *outer-tag-inner-tag*. For IP demux interfaces configured for static subscribers, the VLAN tags configured on the underlying interface are used.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set vlan-tags
```

4. (Optional) Specify that the logical system name is included in the username.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set logical-system-name
```

5. Specify that the routing instance name is included in the username.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set routing-instance-name
```

6. Specify a domain name to include in the username.

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set domain-name domain-name
```

7. (Optional) Specify a delimiter character to separate the username elements except for the domain name. The domain name is always preceded by the @ character. The default delimiter is a period (.)

```
[edit system services static-subscribers group group-name authentication username-include]
user@host# set delimiter delimiter-character
```

Consider the following configuration for the subscriber group, shipping:

```
[edit system services static-subscribers group shipping authentication username-include]
user@host# set user-prefix warehouse3
user@host# set interface
user@host# set logical-system-name
user@host# set routing-instance-name
user@host# set domain-name campus.example.com
```

Configured in the default logical system and routing instance R5 for interface ge-0/1/2.50, this sample configuration generates the following username:

```
warehouse3.ge-0-1-2-50.default.R5@campus.example.com
```

Now consider a different configuration for the same subscriber group, where the static interface has a single-tagged VLAN with an outer VLAN ID of 2101:

```
[edit system services static-subscribers group shipping authentication username-include]
user@host# set user-prefix warehouse3
```

```
user@host# set domain-name Bldg5.example.com
user@host# set vlan-tags
user@host# set delimiter %
```

This sample configuration generates the following username:

warehouse3%2101@Bldg5.example.com

Even though a delimiter of % is configured, the domain name is always separated from preceding elements by @.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.3R1	PS/LT support, introduced in Junos OS Release 18.3R1, enables full subscriber management (equivalent to dynamic subscribers) for statically provisioned subscribers whose traffic is transported over IP/MPLS access models.

RELATED DOCUMENTATION

| [JSRC for Subscriber Provisioning and Accounting](#) | 1856

Monitoring and Management Diameter Information

IN THIS SECTION

- [Verifying Diameter Node, Instance, and Route Information](#) | 1889
- [Verifying and Managing Diameter Application Information](#) | 1890
- [Verifying and Managing Diameter Peer Information](#) | 1892
- [Verifying Diameter Network Element Information](#) | 1894

Verifying Diameter Node, Instance, and Route Information

IN THIS SECTION

- [Purpose | 1889](#)
- [Action | 1889](#)

Purpose

View Diameter node information:

Action

- To display summary information about all Diameter nodes:

```
user@host> show diameter
```

- To display summary information about all Diameter nodes and add information about Diameter applications (functions), instances, network elements, and peers:

```
user@host> show diameter brief
```

- To display brief information about all Diameter nodes and add information about Diameter routes:

```
user@host> show diameter detail
```

- To display summary information about all Diameter instances:

```
user@host> show diameter instance
```

- To display detailed information about all Diameter instances:

```
user@host> show diameter instance detail
```


- To display information about a specific Diameter instance, add the instance name to the command:

```
user@host> show diameter instance master
```

```
user@host> show diameter instance detail master
```

- To display summary information about all Diameter routes:

```
user@host> show diameter route
```

- To display detailed information about all Diameter routes:

```
user@host> show diameter route detail
```

- To display information about a specific Diameter route, add the route name to the command:

```
user@host> show diameter route dne-route2
```

```
user@host> show diameter route detail dne-route2
```

Verifying and Managing Diameter Application Information

IN THIS SECTION

- [Purpose](#) | 1890
- [Action](#) | 1891

Purpose

View or clear Diameter application (function) information:

Action

- To display summary information about all applications associated with Diameter:

```
user@host> show diameter function
```

- To display detailed information about all applications associated with Diameter:

```
user@host> show diameter function detail
```

- To display information about a specific application associated with Diameter, add the application name to the command:

```
user@host> show diameter function jsrc
```

```
user@host> show diameter function detail gx-plus
```

- To display summary statistics about all applications associated with Diameter:

```
user@host> show diameter function statistics
```

- To display detailed statistics about all applications associated with Diameter:

```
user@host> show diameter function statistics detail
```

- To display statistics about a specific application associated with Diameter, add the application name to the command:

```
user@host> show diameter function statistics gx-plus
```

```
user@host> show diameter function statistics detail jsrc
```

- To delete current statistics for all applications associated with Diameter:

```
user@host>clear diameter function statistics
```

- To delete current statistics for a specific application associated with Diameter:

```
user@host>clear diameter function gx-plus statistics
```

Verifying and Managing Diameter Peer Information

IN THIS SECTION

- [Purpose | 1892](#)
- [Action | 1892](#)

Purpose

View or clear Diameter peer information:

Action

- To display summary information about all Diameter peers:

```
user@host> show diameter peer
```

- To display detailed information about all Diameter peers:

```
user@host> show diameter peer detail
```

- To display information about a specific Diameter peer, add the peer name to the command:

```
user@host> show diameter peer peer235
```

```
user@host> show diameter peer detail peer235
```

- To display summary information about Diameter peer-to-network-element mapping for all peers:

```
user@host> show diameter peer map
```

- To display detailed information about Diameter peer-to-network-element mapping for all peers:

```
user@host> show diameter peer map detail
```

- To display information about Diameter peer-to-network-element mapping for a specified peer, add the peer name to the command:

```
user@host> show diameter peer map peer235
```

```
user@host> show diameter peer map detail peer235
```

- To display summary statistics about all Diameter peers:

```
user@host> show diameter peer statistics
```

- To display detailed statistics about all Diameter peers:

```
user@host> show diameter peer statistics detail
```

- To display summary statistics about a specified Diameter peer:

```
user@host> show diameter peer statistics peer235
```

- To display detailed statistics about a specified Diameter peer:

```
user@host> show diameter peer statistics detail peer235
```

- To delete the specified Diameter peer and all of its statistics.

```
user@host>clear diameter peer peer5 connection
```

- To delete the specified Diameter peer and its current statistics:

```
user@host>clear diameter peer peer5 statistics
```

Verifying Diameter Network Element Information

IN THIS SECTION

- [Purpose | 1894](#)
- [Action | 1894](#)

Purpose

View Diameter network element information:

Action

- To display summary information about Diameter network elements:

```
user@host> show diameter network-element
```

- To display detailed information about Diameter network elements:

```
user@host> show diameter network-element detail
```

- To display information about Diameter network elements for a specified network element, include the element name in the command:

```
user@host> show diameter network-element dne-1
```

```
user@host> show diameter network-element detail dne-1
```

- To display summary information about Diameter network-element-to-peer mapping for all network elements:

```
user@host> show diameter network-element map
```

- To display detailed information about Diameter network-element-to-peer mapping for all network elements:

```
user@host> show diameter network-element map detail
```

RELATED DOCUMENTATION

| [Diameter Base Protocol](#) | [1727](#)

Tracing Diameter Base Protocol Events for Troubleshooting

IN THIS SECTION

- [Configuring the Diameter Base Protocol Trace Log Filename](#) | [1896](#)
- [Configuring the Number and Size of Diameter Base Protocol Log Files](#) | [1896](#)
- [Configuring Access to the Diameter Base Protocol Log File](#) | [1897](#)
- [Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged](#) | [1897](#)
- [Configuring the Diameter Base Protocol Tracing Flags](#) | [1898](#)
- [Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged](#) | [1898](#)

The Junos OS trace feature tracks Diameter base protocol 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 `jdiameterd`. 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 Diameter base protocol operations:

Configuring the Diameter Base Protocol Trace Log Filename

By default, the name of the file that records trace output for Diameter base protocol is `jdiameterd`. You can specify a different name with the `file` option.

To configure the filename for Diameter base protocol tracing operations:

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

```
[edit system processes diameter-service traceoptions]
user@host# set file diam_logfile_1
```

Configuring the Number and Size of Diameter Base Protocol 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. (Diameter base protocol supports the files and size options for the traceoptions statement.)

```
[edit system processes diameter-service traceoptions]
user@host# set file diam_1 _logfile_1 files 20 size 2097152
```

Configuring Access to the Diameter Base Protocol 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 diameter-service traceoptions]
user@host# set file diam_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 diameter-service traceoptions]
user@host# set file diam_1 _logfile_1 no-world-readable
```

Configuring a Regular Expression for Diameter Base Protocol 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 diameter-service traceoptions]  
user@host# set file diam_1 _logfile_1 match regex
```

Configuring the Diameter Base Protocol 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 diameter-service traceoptions]  
user@host# set flag dne
```

Configuring the Severity Level to Filter Which Diameter Base Protocol 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. 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* or *verbose*. 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 system processes diameter-service traceoptions]  
user@host# set level severity
```

RELATED DOCUMENTATION

| [Diameter Base Protocol](#) | 1727

Troubleshooting Diameter Networks

IN THIS SECTION

- [Troubleshooting Diameter Network Configuration | 1899](#)
- [Troubleshooting Diameter Network Connectivity | 1900](#)

Troubleshooting Diameter Network Configuration

IN THIS SECTION

- [Problem | 1899](#)
- [Cause | 1900](#)
- [Solution | 1900](#)

Problem

Description

A misconfiguration of the network can prevent Diameter from functioning properly. Configuration options for the Diameter base protocol are simplifying the discovery of a misconfiguration.

Symptoms

The output of the `show diameter peer` command indicates a peer is in the no-activation state. In this case issue the `show diameter peer map` and `show diameter network-element map` commands to determine which network elements use the peer. The output of these commands can indicate why the peer is not activated. For example, all the associated network elements might have higher-priority peers in the open state.

The failed-to-forward counters are increasing in the output of the `show diameter function statistics detail` command. this can indicate that the routes to the peer are incorrectly configured. Check the network connectivity, then use the `show diameter routes` command to determine whether application traffic is being correctly forwarded.

Cause

Typical misconfigurations appear in the routes, peers, and network element configurations.

Solution

Use the appropriate statements to correct the misconfiguration.

Troubleshooting Diameter Network Connectivity

IN THIS SECTION

- [Problem | 1900](#)
- [Cause | 1900](#)
- [Solution | 1901](#)

Problem

Description

In some circumstances, problems can arise with network connectivity to Diameter peers. The problem may originate with the peer or the peer host.

Symptoms

The output of the `show diameter peer` command indicates a peer is in the suspended, rejected, or bad-peer state.

Cause

The suspended state indicates that the peer is not responding or has some other malfunction, or the network path to the peer does not exist.

The rejected state indicates that the network connection has been rejected by the peer.

The bad-peer state indicates that the network connection has been rejected by the router on which the peer resides.

Solution

Determine how persistent the problem is by issuing the `show diameter peer statistics peer-name brief` command to check the connectivity statistics.

RELATED DOCUMENTATION

| [Diameter Base Protocol](#) | [1727](#)

Monitoring and Managing Static Subscriber Information

IN THIS SECTION

- [Forcing a Static Subscriber to Be Logged Out](#) | [1901](#)
- [Resetting the State of an Interface for Static Subscriber Login](#) | [1901](#)
- [Forcing a Group of Static Subscribers to Be Logged Out](#) | [1902](#)
- [Resetting the State of an Interface Group for Static Subscriber Login](#) | [1902](#)
- [Verifying Information about Subscriber Sessions on Static Interfaces](#) | [1902](#)

Forcing a Static Subscriber to Be Logged Out

You can force a static subscriber to be logged out on an interface. After you do so, no subscriber can subsequently log in on that interface until the interface state is reset either by a router reset or by entering the `request services static-subscribers login interface` command.

- To forcibly log out a static subscriber on a static interface:

```
user@host> request services static-subscribers logout interface ge-2/0/1.5
```

Resetting the State of an Interface for Static Subscriber Login

When a static subscriber has been forcibly logged out on an interface with the `request services static-subscribers logout interface` command, you can reset the state of the interface. This action enables a static subscriber to log in on the interface. If you do not reset the state manually, then no static subscribers can log in on the interface until the state is reset by a router reset.

- To reset the state of a static interface:

```
user@host> request services static-subscribers login interface ge-2/0/1.5
```

Forcing a Group of Static Subscribers to Be Logged Out

You can force the static subscribers on all interfaces in a group to be logged out. After you do so, no subscriber can subsequently log in on an interface in that group until the interface state is reset either by a router reset or by entering the `request services static-subscribers login group` command.

- To forcibly log out all static subscribers on a static interface group:

```
user@host> request services static-subscribers logout group boston
```

Resetting the State of an Interface Group for Static Subscriber Login

When static subscribers have been forcibly logged out on an interface group with the `request services static-subscribers logout group` command, you can reset the state of the group. This action enables static subscribers to log in on the interfaces in the group. If you do not reset the state manually, then no static subscribers can log in on any interface in the group until the state is reset by a router reset.

- To reset the state of a static interface group:

```
user@host> request services static-subscribers login group boston
```

Verifying Information about Subscriber Sessions on Static Interfaces

IN THIS SECTION

● [Purpose | 1902](#)

● [Action | 1903](#)

Purpose

View information about subscriber sessions on static interfaces:

Action

- To display information about all static subscriber sessions:

```
user@host> show static-subscribers sessions
```

- To display information about the subscriber sessions for the specified group of static interfaces:

```
user@host> show static-subscribers sessions group boston
```

- To display information about the subscriber session for the specified interface:

```
user@host> show static-subscribers sessions interface ge-0/0/1.1
```

RELATED DOCUMENTATION

| [JSRC and Subscribers on Static Interfaces](#) | 1872

Tracing Static Subscriber Events for Troubleshooting

IN THIS SECTION

- [Configuring the Static Subscribers Trace Log Filename](#) | 1904
- [Configuring the Number and Size of Static Subscribers Log Files](#) | 1904
- [Configuring Access to the Static Subscribers Log File](#) | 1905
- [Configuring a Regular Expression for Static Subscriber Messages to Be Logged](#) | 1905
- [Configuring the Static Subscribers Tracing Flags](#) | 1906
- [Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged](#) | 1906

The Junos OS trace feature tracks static subscriber 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 `jsscd`. 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 static subscriber operations:

Configuring the Static Subscribers Trace Log Filename

By default, the name of the file that records trace output for static subscribers is `jsscd`. You can specify a different name with the `file` option.

To configure the filename for static subscribers tracing operations:

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

```
[edit system processes static-subscribers traceoptions]
user@host# set file stat-subs_1
```

Configuring the Number and Size of Static Subscribers 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 static-subscribers traceoptions]
user@host# set file stat-subs_1 _logfile_1 files 20 size 2097152
```

Configuring Access to the Static Subscribers 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 static-subscribers traceoptions]
user@host# set file stat-subs_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 static-subscribers traceoptions]
user@host# set file stat-subs_1 _logfile_1 no-world-readable
```

Configuring a Regular Expression for Static Subscriber 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 static-subscribers traceoptions]  
user@host# set file stat-subs_1 _logfile match regex
```

Configuring the Static Subscribers 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 static-subscribers traceoptions]  
user@host# set flag authentication
```

Configuring the Severity Level to Filter Which Static Subscriber 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. 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` or `verbose`. 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 system processes static-subscribers traceoptions]  
user@host# set level severity
```

RELATED DOCUMENTATION

| [JSRC and Subscribers on Static Interfaces](#) | 1872

11

PART

Configuration Statements and Operational Commands

- [dynamic-profile \(Domain Map\) | 1908](#)
 - [dynamic-profile \(Static Subscribers\) | 1909](#)
 - [Junos CLI Reference Overview | 1912](#)
-

dynamic-profile (Domain Map)

IN THIS SECTION

- [Syntax | 1908](#)
- [Hierarchy Level | 1908](#)
- [Description | 1908](#)
- [Options | 1908](#)
- [Required Privilege Level | 1909](#)
- [Release Information | 1909](#)

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Dynamic profile that is used for subscriber sessions associated with the domain map.

Options

profile-name—Name of dynamic profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

| [Specifying a Dynamic Profile in a Domain Map](#) | 834

dynamic-profile (Static Subscribers)

IN THIS SECTION

- [Syntax](#) | 1910
- [Hierarchy Level](#) | 1910
- [Description](#) | 1910
- [Default](#) | 1911
- [Options](#) | 1911
- [Required Privilege Level](#) | 1911
- [Release Information](#) | 1911

Syntax

```
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
[edit system services dhcp-local-server dhcpv6]
```

Description

Specify the dynamic client profile that is instantiated at login and de-instantiated at logout for all static subscribers on interfaces configured at the [edit system services static-subscribers interface] hierarchy level or for the static subscribers in a specific group. The group version of the statement takes precedence over the global version.



NOTE: Do not specify a dynamic profile that creates a dynamic interface.

Default

By default, the *junos-default-profile* is used when you do not specify a global dynamic profile with this statement.

Options

profile-name—Name of the dynamic client profile 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

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Subscribers over Static Interfaces Configuration Overview](#) | 1876

[Specifying the Static Subscriber Global Dynamic Profile](#) | 1879

[Specifying the Static Subscriber Group Dynamic Profile](#) | 1884

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)