JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Broadband Subscriber Sessions User Guide

Published
2023-03-26

JUNOS

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

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

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

Junos® OS Broadband Subscriber Sessions User Guide
Copyright © 2023 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 | xlv

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 for Junos OS | **434**
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS | **448**
- DSL Forum Vendor-Specific Attributes | **457**
- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS | **467**
- RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses | **472**
- Support for Cisco Systems VSAs | **473**
- Subscriber Management RADIUS Dictionary Files | **473**
- Interface Text Descriptions for Inclusion in RADIUS Attributes | **473**

RADIUS for Subscriber Management | 476

RADIUS Servers and Parameters for Subscriber Access | **476**

- RADIUS Authentication and Accounting Server Definition | **477**
- Configuring Options that Apply to All RADIUS Servers | **480**
- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | **482**
- Configuring Access Profile Options for Interactions with RADIUS Servers | **483**
- Configuring a Calling-Station-ID with Additional Options | **490**
- Filtering RADIUS Attributes and VSAs from RADIUS Messages | **494**

Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers | **498**

- Interface Description Storage and Reporting Overview | **498**
- Interface Description Storage and Reporting Configuration | **503**

Session Options for Subscriber Access | **503**

- Understanding Session Options for Subscriber Access | **504**
- Subscriber Session Timeout Options | **511**
- Limiting the Number of Active Sessions per Username and Access Profile | **512**
- Configuring Username Modification for Subscriber Sessions | **513**
- Removing Inactive Dynamic Subscriber VLANs | **516**

RADIUS NAS Port Attributes and Options | **518**

- Manual Configuration of the NAS-Port-ID RADIUS Attribute | **518**
- Configuring a NAS-Port-ID with Additional Options | **520**
- Configuring the Order in Which Optional Values Appear in the NAS-Port-ID | **521**
- Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers | **523**
- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview | **524**
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **526**
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **527**
- Manual Configuration of the NAS-Port-Type RADIUS Attribute | **528**
- Configuring the RADIUS NAS-Port-Type per Physical Interface | **531**
- Configuring the RADIUS NAS-Port-Type per VLAN | **532**
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN | **534**
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface | **536**
- Configuring the RADIUS NAS-Port Extended Format per VLAN | **537**
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN | **539**
- Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces | **541**

RADIUS Logical Line Identification | **543**

- RADIUS Logical Line Identifier (LLID) Overview | **544**
- RADIUS Attributes for LLID Preauthentication Requests | **545**
- Configuring Logical Line Identification (LLID) Preauthentication | **546**
- Configuring a Port and Password for LLID Preauthentication Requests | **548**
- Verifying and Managing LLID Preauthentication Configuration | **549**

RADIUS Authentication and Accounting Basic Configuration | **550**

- Configuring Authentication and Accounting Parameters for Subscriber Access | **550**
- Specifying the Authentication and Accounting Methods for Subscriber Access | **551**
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access | **552**

Configuring Local Authentication and Authorization for Subscribers | 552

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 556

Configuring RADIUS Reauthentication for DHCP Subscribers | 567

RADIUS Accounting for Subscriber Access | 570

RADIUS Accounting Statistics for Subscriber Access Overview | 571

RADIUS Acct-On and Acct-Off Messages | 572

Configuring Per-Subscriber Session Accounting | 573

Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI | 576

Understanding RADIUS Accounting Duplicate Reporting | 578

Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting | 580

Configuring Per-Service Session Accounting | 581

Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 583

Configuring Service Packet Counting for Volume Statistics | 585

Configuring Service Accounting | 586

Preservation of RADIUS Accounting Information During an Accounting Server Outage | 588

Configuring Back-up Options for RADIUS Accounting | 591

Forcing the Router to Contact the Accounting Server Immediately | 592

Monitoring Pending RADIUS Accounting Stop Messages | 593

Suspending RADIUS Accounting and Baselining Accounting Statistics Overview | 595

Configuring RADIUS Accounting Suspension and Baselining Accounting Statistics | 599

Verifying and Managing Subscriber AAA Information | 601

Session Termination Causes and RADIUS Termination Cause Codes | 603

Understanding Session Termination Causes and RADIUS Termination Cause Codes | 603

Mapping Session Termination Causes to Custom Termination Cause Codes | 606

AAA Termination Causes and Code Values | 608

DHCP Termination Causes and Code Values | 610

L2TP Termination Causes and Code Values | 611

PPP Termination Causes and Code Values | 638

VLAN Termination Causes and Code Values | 651

Domain Maps for Subscriber Management | 654

Mapping Subscriber Domains to Access and Session Options | 654

Domain Mapping Overview | **655**

Configuring a Domain Map | **659**

Configuring a Wildcard Domain Map | **661**

Specifying an Access Profile in a Domain Map | **662**

Specifying an Address Pool in a Domain Map | **663**

Specifying a Dynamic Profile in a Domain Map | **664**

Specifying an AAA Logical System/Routing Instance in a Domain Map | **664**

Specifying a Target Logical System/Routing Instance in a Domain Map | **665**

Specifying a Tunnel Profile in a Domain Map | **666**

Specifying a Tunnel Switch Profile in a Domain Map | **667**

Configuring Domain and Realm Name Usage for Domain Maps | **667**

Specifying Domain and Realm Name Delimiters | **668**

Specifying the Parsing Order for Domain and Realm Names | **669**

Specifying the Parsing Direction for Domain and Realm Names | **670**

Enabling Domain Name Stripping | **671**

Changing the Username and Password to Simplify Off-Chassis Provisioning | **671**

Verifying Domain Maps | **673**

Testing and Troubleshooting AAA | 675

AAA Testing and Troubleshooting | **675**

AAA Configuration Testing and Troubleshooting | **675**

Testing a Subscriber AAA Configuration | **676**

Tracing General Authentication Service (authd) Events for Troubleshooting | **683**

Configuring the General Authentication Service Trace Log Filename | **684**

Configuring the Number and Size of General Authentication Service Log Files | **684**

Configuring Access to the General Authentication Service Log File | **685**

Configuring a Regular Expression for General Authentication Service Messages to Be Logged | **685**

Configuring Subscriber Filtering for General Authentication Service Tracing | **686**

Configuring the General Authentication Service Tracing Flags | **687**

2 **DHCP and DHCPv6 for Subscriber Management**

DHCP for Subscriber Management | 690

DHCP Overview | **691**

Understanding Differences Between Legacy DHCP and Extended DHCP | **691**

Extended DHCP Relay Agent Overview | **695**

DHCP Relay Proxy Overview | 698

Minimum DHCP Relay Agent Configuration | 700

Example: DHCP Relay Agent Configuration with Multiple Clients and Servers | 700

DHCP Access Profiles for Subscriber Authentication and Accounting Parameters | 702

Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview | 702

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | 703

Attaching an Access Profile to All DHCP Subscriber or All DHCP Client Interfaces | 703

Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients | 704

Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces | 705

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706

Overriding the Default DHCP Local Server Configuration Settings | 707

Overriding the Default DHCP Relay Configuration Settings | 709

DHCP Behavior When Renegotiating While in Bound State | 712

Sending Release Messages When Clients Are Deleted | 713

Disabling Automatic Binding of Stray DHCP Requests | 714

Enabling DHCP Relay Proxy Mode | 716

Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent | 716

Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address | 717

Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall | 717

Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets | 718

Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally | 719

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | 719

Load Balancing DHCP Local Servers by Delaying Responses to Clients | 720

Configuring a Delay in Local Server Response to DHCP Discover and Solicit Messages | 721

DHCP Options and Selective Traffic Processing | 724

DHCP Options and Selective Traffic Processing Overview | 725

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings | 728

Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings | 728

Requirements | 729

Overview | 729

Configuration | 729

Verification | 732

Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing | 734

Requirements | 734

Overview | 735

Configuration | 735

Verification | 738

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | 739

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | 740

Client-Side Support | 742

Server-Side Support | 742

DHCP Local Server Support | 743

DHCP-Initiated Service Change Based on Remote ID | 744

Configuring DHCP-Initiated Service Change Based on Remote ID | 745

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

Using DHCP Option 82 Information | 751

Using DHCP Relay Agent Option 82 Information | 751

Configuring Option 82 Information | 752

Overriding Option 82 Information | 755

Including a Prefix in DHCP Options | 755

Including a Textual Description in DHCP Options | 759

Enable Processing of Untrusted Packets So Option 82 Information Can Be Used | 761

Extracting an Option 82 or Option 37 Substring to Create an Interface Set | 762

Default Services for DHCP Subscribers | 764

Default Subscriber Service Overview | 764

Configuring a Default Subscriber Service | 765

DHCP Client Attribute and Address Assignment | 766

DHCP Attributes Overview | 766

Attributes That Can Be Applied to DHCP Clients | 768

Configuring DHCP Attributes for All Clients or a Group of Clients | 771

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 772

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **773**

Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option | **775**

Specifying the Subnet for DHCP Client Address Assignment | **776**

DHCP Local Server Handling of Client Information Request Messages | **776**

Enabling Processing of Client Information Requests | **777**

DNS Address Assignment Precedence | **778**

Example: Extended DHCP Local Server Configuration with Optional Pool Matching | **779**

DHCP Lease Times for IP Addresses | **780**

DHCP Lease Timers | **780**

DHCP Lease-Time Validation Overview | **781**

Configuring a DHCP Lease-Time Threshold | **783**

DHCP Asymmetric Leasing Overview | **784**

Configuring DHCP Asymmetric Leasing | **786**

DHCP Leasequery Methods | **788**

Benefits of DHCP Leasequery | **789**

DHCP Individual Leasequery | **790**

DHCP Bulk Leasequery | **794**

DHCP Active Leasequery | **799**

Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations | **811**

Configuring and Using DHCP Individual Leasequery | **812**

Configuring and Using DHCP Bulk Leasequery | **814**

Configuring and Using DHCP Active Leasequery | **818**

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | **824**

Verifying and Managing DHCP Individual and Bulk Leasequery Configurations | **829**

Verifying and Managing DHCP Active Leasequery Operations | **830**

DHCP Client Authentication With An External AAA Authentication Service | **833**

Specifying Authentication Support | **833**

Creating Unique Usernames for DHCP Clients | **834**

Example-Configuring DHCP with External Authentication Server | **837**

Receiving DHCP Options From a RADIUS Server | **838**

Centrally Configure DHCP Options on a RADIUS Server | **838**

Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview | **843**

Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6 Subscribers | **846**

Monitoring DHCP Options Configured on RADIUS Servers | **849**

Common DHCP Configuration for Interface Groups and Server Groups | **852**

Grouping Interfaces with Common DHCP Configurations | **852**

Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces | **855**

Configuring Group-Specific DHCP Local Server Options | **856**

Configuring Group-Specific DHCP Relay Options | **857**

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups | **858**

Number of DHCP Clients Per Interface | **862**

Specifying the Maximum Number of DHCP Clients Per Interface | **862**

Allowing Only One DHCP Client Per Interface | **863**

Maintaining DHCP Subscribers During Interface Delete Events | **865**

Maintaining Subscribers During Interface Delete Events | **865**

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | **866**

Configuring an ACX Series DHCP Local Server to Preserve Subscriber Binding Information | **867**

Verifying and Managing DHCP Subscriber Binding During Interface Delete Events | **868**

Dynamic Reconfiguration of Clients From a DHCP Local Server | **870**

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | **870**

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | **874**

Configuring Dynamic Reconfiguration Attempts for DHCP Clients | **875**

Configuring Deletion of the Client When Dynamic Reconfiguration Fails | **876**

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | **876**

Configuring a Token for DHCP Local Server Authentication | **877**

Conserving IP Addresses Using DHCP Auto Logout | **878**

DHCP Auto Logout Overview | **879**

Automatically Logging Out DHCP Clients | **881**

How DHCP Relay Agent Uses Option 82 for Auto Logout | **882**

DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers | **883**

Automatically Logging Out DHCPv6 Clients | **884**

DHCP Short Cycle Protection | **885**

DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions | **885**

Configuring DHCP Short-Cycle Protection | **889**

Verifying and Managing DHCP Short-Cycle Protection | **892**

DHCP Monitoring and Management | **895**

Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings | **895**

Viewing and Clearing DHCP Bindings | **896**

Monitoring DHCP Relay Server Responsiveness | **898**

Verifying DHCP Server Binding and Server Statistics | **899**

Verifying and Managing DHCP Relay Configuration | **901**

Tracing Extended DHCP Operations | **902**

Configuring the Extended DHCP Log Filename | **904**

Configuring the Number and Size of Extended DHCP Log Files | **904**

Configuring Access to the Extended DHCP Log File | **905**

Configuring a Regular Expression for Extended DHCP Messages to Be Logged | **906**

Configuring the Extended DHCP Tracing Flags | **906**

Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged | **907**

Tracing Extended DHCP Operations for Specific Interfaces | **908**

DHCPv6 for Subscriber Management | 910

DHCPv6 Local Server | **910**

DHCPv6 Local Server Overview | **910**

Configuring DHCPv6 Rapid Commit (MX Series, EX Series) | **912**

Preventing Binding of Clients That Do Not Support Reconfigure Messages | **912**

Configuring the DUID Type Supported by DHCPv6 Servers | **913**

Example: Extended DHCPv6 Local Server Configuration | **914**

DHCPv6 Relay Agent | **916**

DHCPv6 Relay Agent Overview | **916**

DHCPv6 Relay Agent Options | **917**

Configuring DHCPv6 Relay Agent Options | **917**

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | **919**

Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets | **921**

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | **923**

DHCPv6 Monitoring and Management | **925**

Requesting DHCPv6 Local Server to Initiate Reconfiguration of Client Bindings | **925**

Verifying and Managing DHCPv6 Local Server Configuration | **927**

Verifying and Managing DHCPv6 Relay Configuration | **928**

3

IPv6 for Subscriber Management

IPv6 for Subscriber Management | 930

Introduction to IPv6 Addresses | **930**

- IPv6 Notation | **931**
- IPv6 Prefixes | **931**
- IPv6 Address Types | **932**

Migration to IPv6 Using IPv4 and IPv6 Dual Stack | **934**

- Basic Architecture of a Subscriber Access Dual-Stack Network | **934**
- Terms Used in IPv6 Subscriber Management Documentation | **935**
- IPv6 Addressing Requirements for a Subscriber Access Network | **937**

IPv6 WAN Link Addressing with NDRA | **939**

- Using NDRA to Provide IPv6 WAN Link Addressing Overview | **939**
- IPv6 Neighbor Discovery Protocol Overview | **941**
- Dynamic Router Advertisement Configuration Overview | **942**
- Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | **942**
- Methods for Obtaining IPv6 Prefixes for NDRA | **944**
- Duplicate Prefix Protection for NDRA | **945**

IPv6 WAN Link Addressing with DHCPv6 IA_NA | **946**

- Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA | **947**
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | **947**

Subscriber LAN Addressing with DHCPv6 Prefix Delegation | **948**

- Using DHCPv6 Prefix Delegation Overview | **949**
- Using a Delegated Prefix on the CPE Loopback Interface | **950**
- DHCPv6 Prefix Delegation over PPPoE | **950**
- Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation | **951**
- DHCPv6 Prefix Exclusion | **952**
- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation | **954**
- Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | **955**

WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **956**

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | **957**
- DHCPv6 Options in a DHCPv6 Multiple Address Environment | **958**

Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA | **959**

Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface | **961**

Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | **961**

Requirements | **961**

Overview | **962**

Configuration | **964**

Verification | **986**

Designs for IPv6 Addressing in a Subscriber Access Network | **993**

Selecting the Type of Addressing Used on the CPE | **993**

Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link | **993**

Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | **994**

Selecting the Method of Obtaining IPv6 Prefixes | **995**

Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **996**

Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | **997**

Design 3: IPv6 Addressing with NDRA | **998**

Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | **999**

Dual-Stack Access Models in a DHCP Network | **1000**

IPv4 and IPv6 Dual Stack in a DHCP Access Network | **1001**

AAA Service Framework in a Dual Stack over a DHCP Access Network | **1002**

Dual-Stack Interface Stack in a DHCP Wholesale Network | **1003**

Single-Session DHCP Dual-Stack Overview | **1004**

Configuring Single-Session DHCP Dual-Stack Support | **1008**

Verifying and Managing DHCP Dual-Stack Configuration | **1011**

Dual-Stack Access Models in a PPPoE Network | **1012**

IPv4 and IPv6 Dual Stack in a PPPoE Access Network | **1013**

Shared IPv4 and IPv6 Service Sessions on PPP Access Networks | **1016**

AAA Service Framework in a Dual Stack over a PPPoE Access Network | **1016**

RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers | **1019**

Accounting Messages for PPPoE Using NDRA Prefixes | **1020**

Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes | **1027**

Suppressing Accounting Information That Comes from AAA | **1037**

Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address | **1038**

Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | **1039**

Best Practice: Static PPPoE Interfaces with NDRA | **1039**

Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network | **1040**

Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | **1041**

Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 | **1041**

Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles | **1042**

Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | **1044**

Dual Stack for PPPoE Access Networks Using DHCP | **1044**

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | **1045**

Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network | **1046**

Dual Stack for PPPoE Access Networks Using NDRA | **1048**

Configuring a Static PPPoE Logical Interface for NDRA | **1049**

Configuring an Address-Assignment Pool Used for Router Advertisements | **1050**

Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | **1051**

Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | **1052**

Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | **1052**

Requirements | **1053**

Overview | **1053**

Configuration | **1055**

Verification | **1073**

Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | **1078**

Requirements | **1078**

Overview | **1079**

Configuration | **1081**

Verification | **1103**

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | **1110**

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | **1111**

Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1114**

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1115**

On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview | **1115**

On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview | **1117**

IPCP Negotiation with Optional Peer IP Address | **1120**

How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled | **1121**

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1122**

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1122**

Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1123**

Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes | **1123**

Enabling IPv4 Release Control VSA (26–164) in RADIUS Messages | **1124**

Dual Stack Subscribers Monitoring and Management | **1125**

Monitoring Active Subscriber Sessions | **1125**

Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance | **1126**

Monitoring Dynamic Subscriber Sessions | **1127**

Monitoring Address Pools Used for Subscribers | **1128**

Monitoring Specific Subscriber Sessions | **1130**

Monitoring the Status of the PPPoE Logical Interface | **1132**

Monitoring Service Sessions for Subscribers | **1133**

Monitoring PPP Options Negotiated with the Remote Peer | **1134**

Monitoring the RADIUS Attribute Used for NDRA | **1135**

4 **Address-Assignment Pools for Subscriber Management**

Address-Assignment Pools for Subscriber Management | 1138

Address-Assignment Pools for Subscriber Management | **1138**

Address-Assignment Pools Overview | **1139**

Address Allocation from Linked Address Pools | **1141**

Address-Assignment Pool Configuration Overview | **1146**

Configuring an Address-Assignment Pool Name and Addresses | **1148**

Configuring a Named Address Range for Dynamic Address Assignment | **1148**

Preventing Addresses from Being Allocated from an Address Pool | **1149**

Configuring Address-Assignment Pool Usage Threshold Traps | **1151**

Configuring Address-Assignment Pool Linking | **1153**

Configuring Address-Assignment Pool Hold-Down | **1154**

Configuring DHCP Local Address Pool Rapid Drain | **1155**

Configuring Static Address Assignment | **1157**

Configuring Duplicate IPv4 Address Protection for AAA | **1158**

Example: Configuring an Address-Assignment Pool | **1160**

Requirements | **1161**

Overview | **1161**

Configuration | **1161**

5 **DNS Addresses for Subscriber Management**

DNS Addresses for Subscriber Management | 1165

DNS Name Server Addresses for Subscriber Management | **1165**

DNS Name Server Address Overview | **1165**

Configuring DNS Name Server Addresses for Subscriber Management | **1167**

Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | **1169**

DNS Resolver for IPv6 DNS Overview | **1170**

Configuring a DNS Server Address for IPv6 Hosts | **1170**

6 **M:N Subscriber Redundancy**

M:N Subscriber Redundancy | 1173

M:N Subscriber Redundancy on BNG | **1173**

M:N Subscriber Redundancy on BNG Overview | **1173**

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization | **1205**

Configure Subscriber Group Redundancy | **1206**

Configure VRRP to Support M:N Redundancy | **1208**

Configure Active Leasequery with Topology Discovery | **1210**

How to Configure M:N Subscriber Redundancy with Pseudowires and DHCP Binding Synchronization | **1211**

Configure Subscriber Group Redundancy | **1212**

Configure Active Leasequery with Topology Discovery | **1215**

Verifying M:N Redundancy and Active Leasequery Topology Discovery Information | **1216**

M:N Subscriber Service Redundancy on DHCP Server | **1219**

M:N Subscriber Service Redundancy on DHCP Server Overview | **1219**

N+1 Support for BNG M:N Subscriber Service Redundancy | **1223**

N+1 Support for BNG M:N Subscriber Service Redundancy Overview | **1224**

How N+1 Support for BNG M:N Subscriber Service Redundancy Works | **1224**

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | **1227**

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Overview | **1227**

How BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Works | **1228**

7

Access Node Control Protocol and the ANCP Agent for Subscriber Services

Access Node Control Protocol and the ANCP Agent for Subscriber Services | 1232

ANCP Agent Neighbors and Operations | **1232**

- ANCP and the ANCP Agent Overview | **1233**
- ANCP Operations in Different Network Configurations | **1243**
- Configuring the ANCP Agent | **1253**
- Configuring ANCP Neighbors | **1255**
- Associating an Access Node with Subscribers for ANCP Agent Operations | **1256**
- Specifying the Interval Between ANCP Adjacency Messages | **1257**
- Specifying the Maximum Number of Discovery Table Entries | **1258**
- Configuring the ANCP Agent for Backward Compatibility | **1258**
- Specifying How Long Processes Wait for the ANCP Agent Restart to Complete | **1259**
- Configuring the ANCP Agent to Learn ANCP Partition IDs | **1260**
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet | **1261**
  - Requirements | **1261**
  - Overview | **1262**
  - Configuration | **1269**
  - Verification | **1287**

ANCP Agent Traffic Shaping and CoS | **1292**

- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Preservation of CoS Shaping Across ANCP Agent Restarts | **1298**
- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | **1304**
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | **1306**
- Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | **1308**
- Verifying and Monitoring CoS for ANCP Subscribers | **1310**

ANCP Agent and AAA | **1311**

- ANCP Agent Interactions with AAA | **1312**
- ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | **1314**
- Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | **1324**

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | **1325**

ANCP Monitoring and Management | **1326**

Triggering ANCP OAM to Test the Local Loop | **1326**

Verifying and Monitoring ANCP Neighbors | **1328**

Clearing ANCP Neighbors | **1329**

Verifying and Monitoring ANCP Subscribers | **1330**

Clearing ANCP Subscribers | **1331**

Clearing and Verifying ANCP Statistics | **1332**

Tracing ANCP Events for Troubleshooting | **1333**

Configuring the ANCP Trace Log Filename | **1334**

Configuring the Number and Size of ANCP Log Files | **1334**

Configuring Access to the ANCP Log File | **1335**

Configuring a Regular Expression for ANCP Messages to Be Logged | **1335**

Configuring the ANCP Tracing Flags | **1336**

Configuring the Severity Level to Filter Which ANCP Messages Are Logged | **1336**

8 **Diameter Base Protocol and its Applications**

Diameter Base Protocol and its Applications | 1338

Diameter Base Protocol | **1338**

Diameter Base Protocol Overview | **1339**

Messages Used by Diameter Applications | **1342**

Diameter AVPs and Diameter Applications | **1350**

Configuring Diameter | **1373**

Configuring the Origin Attributes of the Diameter Instance | **1374**

Configuring Diameter Peers | **1374**

Configuring the Diameter Transport | **1376**

Configuring Diameter Network Elements | **1377**

Example: Configure S6a Application | **1379**

Requirements | **1379**

Overview | **1379**

Configuration | **1380**

Verification | **1389**

Gx-Plus for Provisioning Subscribers | **1392**

Gx-Plus for Provisioning Subscribers Overview | **1393**

Understanding Gx-Plus Interactions Between the Router and the PCRF | **1395**

Configuring Gx-Plus | **1404**

Configuring the Gx-Plus Partition | **1405**

Configuring Gx-Plus Global Attributes | **1406**

Provisioning Subscribers with Gx-Plus | **1407**

Disabling PCRF Control of a Subscriber Session | **1407**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting | **1410**

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | **1418**

Understanding Gy Interactions Between the Router and the OCS | **1431**

Gy File Backup Overview | **1438**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

Understanding Upstream and Downstream Messages for the PCRF | **1444**

Configuring the OCS Partition | **1449**

Configuring the PCRF Partition | **1455**

Configuring OCS Global Parameters | **1462**

NASREQ for Authentication and Authorization | **1463**

Diameter Network Access Server Application (NASREQ) | **1463**

Configuring the Diameter Network Access Server Application (NASREQ) | **1465**

JSRC for Subscriber Provisioning and Accounting | **1467**

Juniper Networks Session and Resource Control (SRC) and JSRC Overview | **1468**

Understanding JSRC-SAE Interactions | **1469**

JSRC Provisioning for Dual-Stack Subscribers | **1472**

JSRC Configuration Overview | **1476**

Configuring the JSRC Partition | **1477**

Assigning a Partition to JSRC | **1478**

Authorizing Subscribers with JSRC | **1478**

Provisioning Subscribers with JSRC | **1478**

Configuring JSRC for Dual-Stack Subscribers | **1479**

Excluding AVPs from Diameter Messages for JSRC | **1480**

Service Accounting with JSRC | **1480**

Configuring Service Accounting with JSRC | **1482**

JSRC and Subscribers on Static Interfaces | **1483**

- Subscribers on Static Interfaces Overview | **1483**
- Subscribers over Static Interfaces Configuration Overview | **1487**
- Example: Configuring Static Subscribers for Subscriber Access | **1488**
- Specifying the Static Subscriber Global Access Profile | **1490**
- Specifying the Static Subscriber Global Dynamic Profile | **1490**
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers | **1491**
- Configuring the Static Subscriber Global Authentication Password | **1492**
- Configuring the Static Subscriber Global Username | **1492**
- Creating a Static Subscriber Group | **1494**
- Specifying the Static Subscriber Group Access Profile | **1495**
- Specifying the Static Subscriber Group Dynamic Profile | **1495**
- Specifying the Static Subscriber Group Service Profile | **1495**
- Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group | **1496**
- Configuring the Static Subscriber Group Authentication Password | **1497**
- Configuring the Static Subscriber Group Username | **1497**

Monitoring and Management Diameter Information | **1499**

- Verifying Diameter Node, Instance, and Route Information | **1500**
- Verifying and Managing Diameter Application Information | **1501**
- Verifying and Managing Diameter Peer Information | **1503**
- Verifying Diameter Network Element Information | **1505**

Tracing Diameter Base Protocol Events for Troubleshooting | **1506**

- Configuring the Diameter Base Protocol Trace Log Filename | **1507**
- Configuring the Number and Size of Diameter Base Protocol Log Files | **1507**
- Configuring Access to the Diameter Base Protocol Log File | **1508**
- Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged | **1508**
- Configuring the Diameter Base Protocol Tracing Flags | **1509**
- Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged | **1509**

Troubleshooting Diameter Networks | **1510**

- Troubleshooting Diameter Network Configuration | **1510**
- Troubleshooting Diameter Network Connectivity | **1511**

Monitoring and Managing Static Subscriber Information | **1512**

- Forcing a Static Subscriber to Be Logged Out | **1512**

Resetting the State of an Interface for Static Subscriber Login | **1512**

Forcing a Group of Static Subscribers to Be Logged Out | **1513**

Resetting the State of an Interface Group for Static Subscriber Login | **1513**

Verifying Information about Subscriber Sessions on Static Interfaces | **1513**

Tracing Static Subscriber Events for Troubleshooting | **1514**

Configuring the Static Subscribers Trace Log Filename | **1515**

Configuring the Number and Size of Static Subscribers Log Files | **1515**

Configuring Access to the Static Subscribers Log File | **1516**

Configuring a Regular Expression for Static Subscriber Messages to Be Logged | **1516**

Configuring the Static Subscribers Tracing Flags | **1517**

Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged | **1517**

9 **Configuration Statements and Operational Commands**

Configuration Statements | 1519

aaa-logical-system (Domain Map) | **1534**

aaa-routing-instance (Domain Map) | **1536**

accept-max-tcp-connections (System Process) | **1537**

accept-sdr (PCRF Partition) | **1539**

access-identifier | **1540**

access-line (Access-Line Rate Adjustment) | **1542**

access-profile | **1557**

access-profile (Extensible Subscriber Services Manager) | **1559**

access-profile (Domain Map) | **1560**

access-profile (Static Subscribers) | **1562**

access-profile-name (Duplicate Accounting) | **1563**

accounting (Access Profile) | **1565**

accounting (Service Accounting) | **1566**

accounting-backup-options (Access Profile) | **1568**

accounting-order (Service Accounting) | **1569**

accounting-stop-on-access-deny | **1571**

accounting-stop-on-failure | **1572**

active-leasequery (DHCP Relay Agent) | **1574**

active-leasequery (DHCP Local Server) | **1577**

active-server-group | **1579**

actual-transit-statistics (Dynamic Profiles) | **1581**

address (Diameter Peer) | **1582**

address (Diameter Transport) | **1584**

address-assignment (Address-Assignment Pools) | **1585**

address-change-immediate-update | **1588**

address-pool (Domain Map) | **1589**

address-protection | **1591**

address-ranges (Demux) | **1593**

adjacency-timer | **1595**

adsl-bytes | **1596**

adsl2-bytes | **1598**

adsl2-plus-bytes | **1600**

advisory-options (Traffic Shaping) | **1602**

aggregate-clients (DHCP Relay Agent) | **1603**

aggregate-clients (Static Subscribers) | **1605**

allow-active-leasequery (DHCP Local Server) | **1608**

allow-bulk-leasequery (DHCP Local Server) | **1610**

allow-leasequery (DHCP Local Server) | **1612**

alternative-partition-name (OCS Partition) | **1614**

always-write-giaddr | **1615**

always-write-option-82 | **1617**

ancp | **1618**

ancp-speed-change-immediate-update (ANCP) | **1621**

asymmetric-lease-time (DHCP Overrides) | **1622**

asymmetric-prefix-lease-time (DHCP Overrides) | **1624**

attempts (DHCP Local Server) | **1626**

attributes (Access-Line Rate Adjustment) | **1628**

attributes (RADIUS Attributes) | **1630**

attributes (JSRC Attributes) | **1633**

authentication (DHCP Local Server) | **1634**

authentication (DHCP Relay Agent) | **1636**

authentication (Static Subscribers) | **1638**

authentication-order | **1640**

authorization-order | **1642**

authentication (Demux) | **1644**

auto-configure (Demux) | **1646**

auto-configure (IPv4) | **1647**

auto-configure (IPv6) | **1650**

autonomous (Dynamic Router Advertisement) | **1652**

backup (OCS Partition) | **1653**

bulk-leasequery (DHCP Relay Agent) | **1655**

called-station-id (OCS Partition) | **1657**

calling-station-id-format (Subscriber Management) | **1659**

charging-id (OCS Partition) | **1661**

charging-service-list | **1662**

circuit-id (DHCP Relay Agent) | **1664**

circuit-type (DHCP Local Server) | **1667**

circuit-type (DHCP Relay Agent) | **1669**

classification-key (DHCP Local Server) | **1671**

classification-key (DHCP Relay Agent) | **1672**

clear-on-abort (DHCP Local Server) | **1675**

client-discover-match (DHCP Local Server) | **1677**

client-discover-match (DHCP Relay Agent) | **1679**

client-id (DHCP Local Server) | **1681**

client-id (DHCP Relay Agent) | **1683**

client-negotiation-match (DHCPv6 Local Server) | **1684**

client-negotiation-match (DHCPv6 Relay Agent) | **1686**

commit-interval | **1687**

coa-immediate-update | **1689**

coa-no-override service-class-attribute | **1690**

concurrent-data-sessions | **1691**

configuration-database (Enhanced Subscriber Management) | **1692**

connect-actively | **1694**

current-hop-limit (Dynamic Router Advertisement) | **1695**

database-replication (Subscriber Session Database) | **1696**

default-action (DHCP Relay Agent Option) | **1698**

default-lifetime (Dynamic Router Advertisement) | **1699**

delay-advertise (DHCPv6) | **1701**

delay-authentication (DHCP Relay Agent) | **1704**

delay-offer (DHCPv4) | **1705**

delegated-pool (DHCP Local Server) | **1708**

delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) | **1710**

delimiter (DHCP Local Server) | **1711**

delimiter (DHCP Relay Agent) | **1714**

delimiter (Domain Map) | **1716**

demux (Interfaces) | **1718**

demux-options (All Demux Interfaces) | **1720**

destination (Diameter Network Element) | **1721**

destination-host | **1723**

destination-host (Gx-Plus) | **1724**

destination-host (OCS Partition) | **1725**

destination-host (PCRF Partition) | **1726**

destination-realm (JSRC) | **1728**

destination-realm (Gx-Plus) | **1729**

destination-realm (OCS Partition) | **1730**

destination-realm (PCRF Partition) | **1732**

dhcp-attributes (Address-Assignment Pools) | **1733**

dhcp-local-server | **1740**

dhcp-relay | **1752**

dhcp-service | **1768**

dhcpv6 (DHCP Local Server) | **1771**

dhcpv6 (DHCP Relay Agent) | **1778**

diameter | **1786**

diameter-instance (JSRC) | **1788**

diameter-instance (Diameter Applications) | **1789**

dictionary | **1791**

disable | **1792**

disable (Extensible Subscriber Services Manager) | **1793**

disable-relay | **1795**

dne-origin (Diameter Network Element) | **1796**

dns-server-address (Dynamic Profiles) | **1798**

domain (Domain Map) | **1800**

domain-name (DHCP Local Server) | **1801**

domain-name (DHCP Relay Agent) | **1804**

domain-name (Static Subscribers) | **1806**

domain-name-server (Routing Instances and Access Profiles) | **1807**

domain-name-server-inet (Routing Instances and Access Profiles) | **1809**

domain-name-server-inet6 (Routing Instances and Access Profiles) | **1811**

downstream-rate (Traffic Shaping) | **1812**

draining (Diameter Applications) | **1814**

draining-response-timeout (Diameter Applications) | **1815**

drop (DHCP Relay Agent Option) | **1817**

dsl (Access-Line Rate Adjustment) | **1818**

dual-stack (DHCP Local Server Overrides) | **1824**

dual-stack (DHCP Relay Agent Overrides) | **1825**

dual-stack-group (DHCP Local Server) | **1827**

dual-stack-group (DHCP Relay Agent) | **1830**

dual-stack-interface-client-limit (DHCP Local Server and Relay Agent) | **1833**

dualstack-support (JSRC) | **1834**

duplication (Access Profile) | **1836**

duplication-filter (Access Profile) | **1837**

duplication-vrf (Duplicate Accounting) | **1839**

dynamic-profile (Demux) | **1840**

dynamic-profile (DHCP Local Server) | **1842**

dynamic-profile (DHCP Relay Agent) | **1844**

dynamic-profile (Domain Map) | **1846**

dynamic-profile (Static Subscribers) | **1847**

dynamic-profiles | **1849**

enable | **1863**

enable (Enhanced Subscriber Management) | **1864**

equals (DHCP Relay Agent) | **1865**

exceed-action | **1869**

exclude (JSRC Attributes) | **1871**

exclude (RADIUS Attributes) | **1872**

excluded-address (Address-Assignment Pools) | **1881**

excluded-range (Address-Assignment Pools) | **1882**

external-authority | **1884**

failover (System Process) | **1885**

family (Address-Assignment Pools) | **1886**

family-state-change-immediate-update | **1888**

final-response-timeout (OCS Partition) | **1890**

force-continue (OCS Partition) | **1891**

forward-only (DHCP Relay Agent Option) | **1893**

forward-only (DHCP Relay Agent) | **1894**

forward-only-replies (DHCP Relay Agent) | **1897**

forwarding (Diameter Network Element) | **1898**

function (Diameter Network Element) | **1899**

function (Diameter Route) | **1901**

ggsn-address (OCS Partition) | **1903**

ggsn-mcc-mnc (OCS Partition) | **1904**

global (Gx-Plus) | **1906**

global (OCS) | **1907**

global (PCRF) | **1908**

group (DHCP Local Server) | **1910**

group (DHCP Relay Agent) | **1915**

group (Static Subscribers) | **1921**

gsmp-syn-timeout (ANCP) | **1923**

gsmp-syn-wait (ANCP) | **1924**

gx-plus (Gx-Plus) | **1926**

host (Address-Assignment Pools) | **1927**

host-name (DHCP Relay Agent) | **1929**

host-name (DHCPv6 Relay Agent) | **1930**

ietf-mode | **1931**

immediate-update | **1933**

include-ipv6 (Gx-Plus) | **1934**

include-irb-and-l2 | **1935**

include-option-82 (DHCP Local Server) | **1938**

inet (Interfaces) | **1940**

inet6 (Interfaces) | **1942**

interface (DHCP Local Server) | **1944**

interface (DHCP Relay Agent) | **1947**

interface (Dynamic Router Advertisement) | **1950**

interface (Static Subscriber Group) | **1952**

interface (Static Subscriber Username) | **1954**

interface-client-limit (DHCP Local Server) | **1955**

interface-client-limit (DHCP Relay Agent) | **1958**

interface-delete (Subscriber Management or DHCP Client Management) | **1960**

interface-description (DHCP Local Server) | **1961**

interface-description (DHCP Relay Agent) | **1963**

interface-description-format | **1965**

interface-name (DHCP Local Server) | **1967**

interface-name (DHCP Relay Agent) | **1968**

interface-mib (Enhanced Subscriber Management) | **1970**

interface-set (ANCP) | **1971**

interface-traceoptions (DHCP) | **1973**

interfaces (ANCP) | **1975**

interfaces (Static and Dynamic Subscribers) | **1977**

interim-rate (Access) | **1984**

ip-address-first | **1985**

ip-can-type (PCRF Partition) | **1986**

jsrc (JSRC) | **1988**

jsrc (Access Profile) | **1989**

jsrc-partition | **1991**

layer2-unicast-replies | **1992**

keep-incoming-circuit-id (DHCP Relay Agent) | **1993**

keep-incoming-interface-id (DHCP Relay Agent) | **1995**

keep-incoming-remote-id (DHCP Relay Agent) | **1996**

leasequery (DHCP Relay Agent) | **1998**

lease-time-threshold (DHCP Local Server and DHCP Relay Agent) | **2000**

lease-time-validation (DHCP Local Server and DHCP Relay Agent) | **2002**

limit | **2003**

linked-pool-aggregation (Address-Assignment Pools) | **2004**

local (Flat-File Access Profile) | **2006**

local-decision (PCRF Partition) | **2008**

local-server-group (DHCP Relay Agent Option) | **2011**

location (DHCP Relay Agent) | **2013**

location (DHCPv6 Relay Agent) | **2014**

logical-interface-unit-range | **2015**

logical-system (Diameter Peer) | **2017**

logical-system (Diameter Transport) | **2018**

logical-system-name (Static Subscribers) | **2020**

logical-system-name (DHCP Local Server) | **2021**

logical-system-name (DHCP Relay Agent) | **2023**

logout-response-timeout (PCRF Partition) | **2025**

ltv-syslog-interval (System Process) | **2026**

mac-address (DHCP Local Server) | **2028**

mac-address (DHCP Relay Agent) | **2030**

maintain-subscriber (Subscriber Management) | **2031**

managed-configuration (Dynamic Router Advertisement) | **2033**

map (Domain Map) | **2034**

max-advertisement-interval (Dynamic Router Advertisement) | **2037**

max-data-sessions-per-subscriber | **2039**

max-db-size (Enhanced Subscriber Management) | **2040**

max-failures | **2043**

max-outstanding-requests (Diameter Applications) | **2044**

max-pending-accounting-stops (Access Profile) | **2046**

max-withhold-time (Access Profile) | **2047**

maximum-discovery-table-entries | **2049**

maximum-helper-restart-time | **2050**

maximum-subscribers | **2052**

metric (Diameter Route) | **2053**

min-advertisement-interval (Dynamic Router Advertisement) | **2054**

multi-address-embedded-option-response (DHCP Local Server) | **2056**

nas-port-extended-format | **2057**

nas-port-extended-format (Interfaces) | **2060**

nas-port-id-format (Subscriber Management) | **2062**

nas-port-options (RADIUS Options) | **2065**

nas-port-type (Subscriber Management) | **2067**

nas-port-type (RADIUS Options) | **2069**

nasreq (Diameter Application) | **2072**

neighbor (Define ANCP) | **2074**

network | **2076**

network-element (Diameter Base Protocol) | **2077**

network-services | **2079**

no-bind-on-request (DHCP Relay Agent) | **2081**

no-unsolicited-ra (Enhanced Subscriber Management) | **2083**

no-vlan-interface-name | **2084**

not-present (DHCP Relay Agent) | **2087**

ocs (Diameter Applications) | **2090**

on-demand-ip-address | **2093**

on-demand-address-allocation | **2095**

on-link (Dynamic Router Advertisement) | **2096**

option-order (DHCP Relay Agent) | **2098**

option-15 (DHCP Relay Agent) | **2100**

option-16 (DHCP Relay Agent) | **2103**

option-60 (DHCP Local Server) | **2105**

option-60 (DHCP Relay Agent) | **2107**

option-77 (DHCP Relay Agent) | **2110**

option-82 (DHCP Relay Agent) | **2112**

option-82 (DHCP Local Server Authentication) | **2114**

option-82 (DHCP Local Server Pool Matching) | **2116**

option-82 (Address-Assignment Pools) | **2117**

option-match | **2119**

option-number (DHCP Relay Agent Option) | **2121**

options (Access Profile) | **2122**

order | **2132**

origin (Diameter Base Protocol) | **2134**

other-bytes | **2135**

other-overhead-adjust | **2137**

other-stateful-configuration (Dynamic Router Advertisement) | **2139**

overhead-accounting (ANCP) | **2140**

override-chap-password | **2142**

override-password (Domain Map) | **2143**

overrides (DHCP Local Server) | **2144**

overrides (DHCP Relay Agent) | **2148**

overrides (Enhanced Subscriber Management) | **2150**

parse-direction (Domain Map) | **2154**

parse-order (Domain Map) | **2155**

partition | **2157**

partition (Gx-Plus) | **2158**

partition (NASREQ Diameter Application) | **2159**

partition (OCS) | **2161**

partition (PCRF) | **2163**

partition (s6a) | **2166**

password (Static Subscribers) | **2168**

password (DHCP Local Server) | **2170**

password (DHCP Relay Agent) | **2172**

pcrf (Diameter Applications) | **2174**

peer (Diameter Base Protocol) | **2176**

peer (Diameter Network Element) | **2178**

peer-ip-address-optional | **2179**

peer-origin (Diameter Peer) | **2181**

pon (Access-Line Rate Adjustment) | **2182**

pool (Address-Assignment Pools) | **2187**

pool (DHCP Local Server Overrides) | **2190**

pool-match-order | **2192**

port (Diameter Peer) | **2194**

pre-ietf-mode | **2195**

preauthentication-order (Access Profile) | **2196**

preferred-lifetime (Dynamic Router Advertisement) | **2198**

prefix (DHCP Relay Agent) | **2199**

prefix (Address-Assignment Pools) | **2201**

prefix (Dynamic Router Advertisement) | **2203**

priority (Diameter Peer) | **2204**

profile (Access) | **2205**

process-inform | **2213**

protocol-master | **2215**

protocols (Dynamic Profiles) | **2218**

provisioning-order (Diameter Applications) | **2221**

proxy-mode | **2223**

qos-adjust | **2225**

qos-adjust-adsl | **2227**

qos-adjust-adsl2 | **2228**

qos-adjust-adsl2-plus | **2230**

qos-adjust-other | **2232**

qos-adjust-sdsl | **2234**

qos-adjust-vdsl | **2236**

qos-adjust-vdsl2 | **2238**

radius (Access Profile) | **2239**

radius-disconnect (DHCP Local Server) | **2244**

radius-flow-tap | **2246**

radius-options (Access) | **2250**

radius-options (Interfaces) | **2251**

radius-server | **2253**

range (Address-Assignment Pools) | **2259**

rapid-commit (DHCPv6 Local Server) | **2261**

reachable-time (Dynamic Router Advertisement) | **2262**

realm-delimiter (Domain Map) | **2264**

realm-parse-direction (Domain Map) | **2265**

reauthenticate (DHCP Local Server) | **2267**

reconfigure (DHCP Local Server) | **2270**

redundancy (M:N Subscriber Redundancy) | **2272**

relay-agent-interface-id (DHCP Local Server) | **2276**

relay-agent-interface-id (DHCPv6 Relay Agent) | **2278**

relay-agent-interface-id (DHCPv6 Relay Agent Username) | **2280**

relay-agent-remote-id (DHCP Local Server) | **2281**

relay-agent-remote-id (DHCPv6 Relay Agent) | **2283**

relay-agent-remote-id (DHCPv6 Relay Agent Username) | **2285**

relay-agent-subscriber-id (DHCP Local Server) | **2287**

relay-agent-subscriber-id (DHCPv6 Relay Agent) | **2288**

relay-option (DHCP Relay Agent) | **2290**

relay-option-vendor-specific (dhcpv6) | **2292**

relay-option-82 | **2293**

relay-server-group (DHCP Relay Agent Option) | **2296**

relay-source | **2298**

remote-id (DHCP Relay Agent) | **2300**

remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) | **2303**

replace-ip-source-with (DHCP Relay Agent) | **2305**

report-interface-descriptions (Access) | **2307**

report-local-rule (PCRF Partition) | **2308**

report-resource-allocation (PCRF Partition) | **2310**

report-successful-resource-allocation (PCRF Partition) | **2312**

request-max-tcp-connections (System Process) | **2314**

request-rate (Access) | **2315**

requested-ip-network-match (DHCP Local Server) | **2317**

retransmit-timer (Dynamic Router Advertisement) | **2318**

revert-interval (Access) | **2320**

route (Diameter Network Element) | **2321**

router-advertisement (Dynamic Profiles) | **2323**

routing-instance (Diameter Peer) | **2324**

routing-instance (Diameter Transport) | **2326**

routing-instance-name (DHCP Local Server) | **2327**

routing-instance-name (DHCP Relay Agent) | **2329**

routing-instance-name (Static Subscribers) | **2331**

s6a | **2333**

sdsl-bytes | **2335**

sdsl-overhead-adjust | **2336**

send-acct-status-on-config-change (Access Profile) | **2338**

send-dyn-subscription-indicator (PCRF Partition) | **2340**

send-network-family-indicator (PCRF Partition) | **2342**

send-origin-state-id (Diameter Applications) | **2344**

send-release-on-delete (DHCP Relay Agent) | **2345**

server-duid-type (DHCP Local Server) | **2347**

server-group | **2348**

server-id-override | **2350**

server-response-time (DHCP Relay Agent) | **2352**

service (Service Accounting) | **2354**

service-context-id (OCS) | **2355**

service-profile (DHCP Local Server) | **2357**

service-profile (DHCP Relay Agent) | **2359**

service-profile (Static Subscribers) | **2361**

services (System Services) | **2362**

session-limit-per-username (Access Profile) | **2370**

session-options | **2372**

sftp-backup (OCS Partition) | **2376**

shmlog (Shared Memory Log) | **2378**

short-cycle-protection (DHCP Local Server and Relay Agent) | **2382**

smg-service (Enhanced Subscriber Management) | **2384**

source-interface-set-at-login | **2385**

stacked-vlan-ranges (RADIUS Options) | **2387**

starts-with (DHCP Relay Agent Option) | **2389**

static-subscribers (Dynamic Service Provisioning) | **2393**

statistics (Access Profile) | **2395**

statistics (Service Accounting) | **2397**

strict (DHCP Local Server) | **2398**

strip-domain (Domain Map) | **2400**

strip-username (Domain Map) | **2401**

sub-domain | **2402**

subscriber (Access Profile) | **2408**

subscriber-packet-idle-timeout | **2410**

subscriber-management (Subscriber Management) | **2412**

subscriber-profile | **2414**

subscription-id-data-include (PCRF Partition) | **2416**

subscription-id-type (PCRF Partition) | **2418**

target-logical-system (Domain Map) | **2420**

target-routing-instance (Domain Map) | **2422**

terminate-code | **2424**

timeout (DHCP Local Server) | **2426**

timeout-grace (Access) | **2428**

token (DHCP Local Server) | **2430**

trace (DHCP Local Server) | **2432**

trace (DHCP Relay Agent) | **2433**

traceoptions (ANCP) | **2435**

traceoptions (DHCP) | **2438**

traceoptions (Diameter Base Protocol) | **2441**

traceoptions (Extensible Subscriber Services Manager) | **2444**

traceoptions (Enhanced Subscriber Management) | **2445**

traceoptions (General Authentication Service) | **2449**

traceoptions (Static Subscribers) | **2451**

transport (Diameter Base Protocol) | **2454**

transport (Diameter Peer) | **2455**

traps | **2457**

trigger (DHCP Local Server) | **2459**

trio-flow-offload | **2461**

trust-option-82 | **2462**

tunnel-profile (Domain Map) | **2464**

underlying-interface (ANCP) | **2465**

unique-nas-port (Access) | **2466**

unit | **2468**

unit (Dynamic Profiles Standard Interface) | **2479**

update-interval | **2483**

update-interval (Service Accounting) | **2485**

update-response-timeout (PCRF Partition) | **2487**

upstream-rate (Traffic Shaping) | **2489**

use-interface-description | **2490**

username-include (Demux) | **2493**

username-include (DHCP Local Server) | **2495**

username-include (DHCP Relay Agent) | **2497**

user-name-include (OCS Partition) | **2500**

username-include (Static Subscribers) | **2503**

use-option-82 | **2505**

use-primary (DHCP Relay Agent) | **2507**

use-underlying-interface-mac | **2509**

use-vlan-id | **2510**

use-vlan-id (DHCP Relay Agent) | **2512**

user-prefix (DHCP Local Server) | **2514**

user-prefix (DHCP Relay Agent) | **2516**

user-prefix (Static Subscribers) | **2518**

valid-lifetime (Dynamic Router Advertisement) | **2520**

vdsl-bytes | **2521**

vdsl-overhead-adjust | **2523**

vdsl2-bytes | **2525**

vdsl2-overhead-adjust | **2526**

vendor-specific (DHCP Relay Agent) | **2528**

violation-action (DHCP Local Server and DHCP Relay Agent) | **2530**

vlan-ranges (RADIUS Options) | **2531**

vrf-name (Duplicate Accounting) | **2533**

wait-for-acct-on-ack (Access Profile) | **2535**

Operational Commands | 2537

clear ancp neighbor | **2541**

clear ancp statistics | **2543**

clear ancp subscriber | **2547**

clear dhcp relay active-leasequery statistics | **2549**

clear dhcp relay binding | **2553**

clear dhcp relay lockout-entries | **2556**

clear dhcp relay statistics | **2559**

clear dhcp server active-leasequery statistics | **2562**

clear dhcp server binding | **2563**

clear dhcp server lockout-entries | **2568**

clear dhcp server statistics | **2570**

clear dhcpv6 relay active-leasequery statistics | **2573**

clear dhcpv6 relay binding | **2576**

clear dhcpv6 relay lockout-entries | **2580**

clear dhcpv6 relay statistics | **2583**

clear dhcpv6 server active-leasequery statistics | **2585**

clear dhcpv6 server binding | **2587**

clear dhcpv6 server lockout-entries | **2591**

clear dhcpv6 server statistics | **2593**

clear diameter function statistics | **2595**

clear diameter peer | **2597**

clear extensible-subscriber-services counters | **2598**

clear extensible-subscriber-services sessions | **2599**

clear ipv6 router-advertisement | **2601**

clear network-access aaa statistics | **2602**

clear network-access aaa subscriber | **2607**

clear network-access gx-plus replay | **2610**

clear network-access gx-plus statistics | **2611**

clear network-access ocs statistics | **2613**

clear network-access pcrf | **2614**

clear system subscriber-management statistics | **2616**

request ancp oam interface | **2617**

request ancp oam neighbor | **2619**

request ancp oam port-down | **2621**

request ancp oam port-up | **2623**

request dhcp relay bulk-leasequery | **2625**

request dhcp relay leasequery | **2628**

request dhcp server reconfigure | **2630**

request dhcpv6 server reconfigure | **2632**

request dhcpv6 relay bulk-leasequery | **2634**

request dhcpv6 relay leasequery | **2637**

request network-access aaa accounting | **2639**

request network-access aaa replay pending-accounting-stops | **2641**

request network-access aaa subscriber modify session-id | **2642**

request network-access aaa subscriber set session-id | **2644**

request services extensible-subscriber-services reload-dictionary | **2646**

request services static-subscribers login group | **2648**

request services static-subscribers logout group | **2649**

request services static-subscribers login interface | **2651**

request services static-subscribers logout interface | **2652**

request services subscribers | **2654**

request services subscribers clear | **2655**

request system reboot (Junos OS) | **2657**

restart extensible-subscriber-services | **2667**

show accounting pending-accounting-stops | **2669**

show ancp cos | **2674**

show ancp neighbor | **2681**

show ancp statistics | **2693**

show ancp subscriber | **2699**

show ancp summary | **2709**

show ancp summary neighbor | **2712**

show ancp summary subscriber | **2715**

show class-of-service interface | **2717**

show class-of-service interface-set | **2762**

show class-of-service scheduler-map | **2765**

show class-of-service traffic-control-profile | **2770**

show database-replication statistics | **2776**

show database-replication summary | **2778**

show dhcp relay active-leasequery | **2782**

show dhcp relay binding | **2789**

show dhcp relay lockout-entries | **2796**

show dhcp relay statistics | **2800**

show dhcp server active-leasequery | **2807**

show dhcp server active-leasequery statistics | **2809**

show dhcp server binding | **2811**

show dhcp server lockout-entries | **2821**

show dhcp server statistics | **2825**

show dhcpv6 relay active-leasequery | **2831**

show dhcpv6 relay binding | **2838**

show dhcpv6 relay lockout-entries | **2850**

show dhcpv6 relay statistics | **2854**

show dhcpv6 server active-leasequery | **2860**

show dhcpv6 server active-leasequery statistics | **2862**

show dhcpv6 server binding | **2865**

show dhcpv6 server lockout-entries | **2874**

show dhcpv6 server statistics | **2877**

show diameter | **2883**

show diameter function | **2892**

show diameter function statistics | **2898**

show diameter instance | **2903**

show diameter network-element | **2905**

show diameter network-element map | **2910**

show diameter peer | **2914**

show diameter peer map | **2920**

show diameter peer statistics | **2924**

show diameter route | **2929**

show dynamic-profile session | **2932**

show ipv6 router-advertisement | **2939**

show network-access aaa accounting | **2946**

show network-access aaa radius-servers | **2948**

show network-access aaa statistics | **2958**

show network-access aaa statistics authentication | **2974**

show network-access aaa statistics pending-accounting-stops | **2979**

show network-access aaa statistics preauthentication | **2980**

show network-access aaa statistics re-authentication | **2983**

show network-access aaa subscribers | **2985**

show network-access aaa subscribers session-id | **2992**

show network-access aaa terminate-code | **3003**

show network-access address-assignment pool | **3010**

show network-access domain-map | **3012**

show network-access gx-plus | **3014**

show network-access nasreq statistics | **3032**

show network-access ocs | **3037**

show network-access pcrf | **3040**

show network-access requests statistics | **3045**

show network-access s6a | **3048**

show ppp address-pool | **3051**

show static-subscribers sessions | **3053**

show subscribers | **3057**

show subscribers summary | **3110**

show system subscriber-management redundancy-state dhcp active-leasequery interface | **3121**

show system subscriber-management redundancy-state interface | **3124**

show system subscriber-management route | **3127**

show system subscriber-management statistics | **3134**

show system subscriber-management summary | **3147**

test aaa authd-lite user | **3151**

test aaa dhcp user | **3156**

test aaa ppp user | **3164**

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.

PART 1

AAA for Subscriber Management

AAA for Subscriber Management | 2

RADIUS for Subscriber Management | 476

Domain Maps for Subscriber Management | 654

Testing and Troubleshooting AAA | 675

CHAPTER 1

AAA for Subscriber Management

IN THIS CHAPTER

- AAA Service Framework Overview | 2
- Standard and Vendor-Specific RADIUS Attributes | 3

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 | **483**

Standard and Vendor-Specific RADIUS Attributes | **3**

RADIUS Authentication and Accounting Basic Configuration | **550**

RADIUS Accounting for Subscriber Access | **570**

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 for Junos OS | **434**
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS | **448**
- DSL Forum Vendor-Specific Attributes | **457**
- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS | **467**
- RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses | **472**
- Support for Cisco Systems VSAs | **473**
- Subscriber Management RADIUS Dictionary Files | **473**

- Interface Text Descriptions for Inclusion in RADIUS Attributes | **473**

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 *Junos OS 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 1: Supported RADIUS IETF Attributes

Attribute Number	Attribute Name	Description	Dynamic CoA Support
1	User-Name	• Name of user to be authenticated. • Configurable username override. • Non-standard use for LLID preauthentication feature.	No
2	User-Password	• 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	Value provided by a PPP (CHAP) user in response to the challenge. 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.	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	• 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	• 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	Name of a subscriber firewall filter, formatted as follows: • For an IPv4 input filter—IPv4-ingress:*ingress-filter-name* • For an IPv4 output filter—IPv4-egress:*egress-filter-name* • For an IPv6 input filter—IPv6-ingress:*ingress-filter-name* • For an IPv6 output filter—IPv6-egress:*egress-filter-name* RADIUS accounting request messages, Acct-Start and Acct-Stop, can include more than one Filter-Id attribute, one of each of the listed types. However, RADIUS Access-Accept messages can include only one attribute instance. The value is always treated as an IPv4 input filter name.	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	• 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	String that provides routing information to be configured for the user on the NAS in the format: `<addr>[/<maskLen>] [<nexthop> [<cost>]] [tag <tagValue>] [distance <distValue>]` 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. 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.	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	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: • `decimal`—For example, `435264` • `description`—In the generic format, `jnpr` *`interface-specifier`*`:`*`subscriber-session-id`*; For example, `jnpr fastEthernet 3/2.6:1010101010101`	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	Reason the service (a PPP session) was terminated. The service can be terminated for the following reasons: • 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	• 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	• 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	• 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	• 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	Number of seconds between each interim accounting update for this session. The router uses the following guidelines for interim accounting: • 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). **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).	No

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

Attribute Number	Attribute Name	Description	Dynamic CoA Support
87	NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user. 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: *media:local address:peer address:local tunnel id:peer tunnel id:local session id:peer session id:call serial number*. 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.	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	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. • 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
		• 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 OS AAA Service Framework. The AAA Service Framework uses vendor ID 4874, which is assigned to Juniper Networks by the Internet Assigned Numbers Authority (IANA). Some VSAs correspond to Juniper Networks predefined variables; see *Junos OS 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 logical system:routing instance name. Allowed	string: *logical system:routing instance*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		only from AAA server for default logical system:rout in		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ginstance. When this VSA is not included in the subs		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		cri be r pr of il e, th e r ou ti n 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
		ubscriber_theoneinwhichthesubscribersession		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		comesup—variesbysubscribertype. For DHCPa		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ndPPPoEsubscribers, it is the default routing instan		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ce. For L2TP tunnel subscribers, it is the routing inst		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rnon-default. If the tunnel routing instance is not de		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		erVSAtosetthedesiredroutinginstance.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-4	Primary-DNS	Client DNS address negotiated during IPCP.	integer: 4-byte *primary-dns-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-5	Secondary-DNS	Client DNS address negotiated during IPCP	integer: 4-byte *secondary-dns-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		PCP.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-7	Secondary-WINS	Client WINS (NBNS) address negotiated during IPCP	integer: 4-byte *secondary-wins-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		PCP.		
26-8	Tunnel-Virtual-Router	Virtual router name for tunnel connection.	string: *tunnel-virtual-router*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-9	Tunnel-Password	Tunnel password in clear text. Do not use both ht	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 because set the password is encry		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		pted 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 policy name to apply to client interface.	string: *output-policy-name*	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-23	IGMP-Enable	Whether IGMP is enabled or disabled on a client int	integer: • 0 = disable • 1 = enable	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		erface.		
26-24	PPPoE-Description	ClientMACaddress.	string:pppoe *client-mac-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-25	Redirect-VRouter-Name	Client logical system:routing instance name indicating	string: *logical-system:routing-instance*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		to which logical system:routing instance the request is		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-30	Tunnel-Nas-Port-Method	Method that determines whether the RADIUS	4-octet integer: • 0 = none • 1 = CiscoCLID	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		server convey stothe LNS thephysical NAS port		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		number identifier and the type of the physical port		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		veyedonlywhenthe VSA value is 1. The VSA is form		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		attedsuchthatthefirstoctetindicatesthetunn		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		el and the remaining three bytes are the attribute		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		value.		
26-31	Service-Bundle	SSC service bundle.	string *bundle-name*	No

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	Route tag to apply to or return framed-ip-add	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ress.		

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 attribute rolls o	integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		verits4-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 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 primary IPv6 DNS address negotiated by	hexadecimal string: *ipv6-primary-dns-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		DHCP.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-48	Ipv6-Secondary-DNS	Client secondary IPv6 DNS address negotiate	hexadecimal string: *ipv6-secondary-dns-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-51	Disconnect-Cause	Disconnect cause when a tunneled subscriber is di	hexadecimal string: *disconnect-cause*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		s c o n n e c t e d , a n d L 2 T P l a y e r o f t h e L N S i n i t i a t e s t h		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		termination. The PPP Disconnect 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
		topmessagethatthe routerssendstothe RAD		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		IUSserver.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-55	DHCP-Options	Client DHCP options. Starting in Junos OS Release	hexadecimal string: *dhcp-options*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		17.4R1, include esonly DHCPv4 options. Inearlierreleases		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		es, include sboth DHCPv4 and DHCPv6 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	Traffic mirroring action. For dynamic CoA, VSA 26-5	salt-encrypted integer 0 = stop mirroring 1 = start mirroring	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8 change sthe action on the mirrored traffic identified	2 = no action	

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		deany of the RADIUS-based mirroring attributes		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		include all four VSAs. If the CoA action is to stop mirr		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		other three attributes in the CoA message must

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		atch 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	Identifier that associates mirrored traffic to a spe	salt-encrypted string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		cific subscriber. For dynamic CoA, VSA 26-58 chan		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gesthe action on the mirrored 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 can include any		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		of the RADIUS-based mirroring 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 VSAs.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-60	Med-Ip-Address	IP address of content destination device to whi	salt-encrypted IP address	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		chmirroredtraffic isforwarded. CoA-Requestmes		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sagesthatincludeanyoftheRADIUS-basedmir		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		roring attributes (VSAs 26-58, 26-59, 26-60, or 26		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

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 content destination device to which	salt-encrypted integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hmirroredtraffic is forwarded. CoA-Request message		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		agesthatincludeanyofthe RADIUS-basedmirr		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		oring attributes (VSAs 26-58, 26-59, 26-60, or 26-		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-63	Interface-Desc	Text string that identifies the subscriber's access	string: *interface-description*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sinterface.		

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 add	string: *tunnel-group-name*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		omainmap.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-65	Activate-Service	Service to 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 to deactivate for the subscriber.	string: *service-name*	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-67	Service-Volume	Amount of traffic, in MB, that can use these services;	integer • range=0 through 16777215 MB • 0=no	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		service 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 that these services can be active	integer • range=0 through 16777215 seconds	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ive; service is deactivated when the timeout expire	• 0 = no timeout	

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 statistics for this service is enabled or di	integer • 0 = disable • 1 = enable time statistics	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sabled.TaggedVSA, whichsupports8tags(1-	•	2=enableandvolumestatistics

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-octet	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ter.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-77	MLD-Version	MLD protocol version.	integer: 1-octet • 1 = MLD version 1 • 2 = MLD versi	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				on2

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-78	IGMP-Version	IGMP protocol version.	integer: 1-octet • 1 = IGMPversion1 • 2 = IGMPver	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			• session2 • 3=IGMPversion3	

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
26-91	Tunnel-Switch-Profile	Tunnel switch profile that determines whether	string: *profile-name*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		asubscribersessionisswitchedtoasecondsessi		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ontoaremoteLNS.Takesprecedenceoverun		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		nelswitchprofilesappliedinanyothermanner.		

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 access loop parameter (A	string: actual upstream rate access loop parameter (ASCII enc	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		SCII encoded) as defined in GSMP extensions for La	oded)	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		yer2control(L2C)TopologyDiscoveryandLineCo		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		nfiguration.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-93	L2C-Down-Stream-Data	Actual downstream rate access loop parameter	string: actual downstream rate access loop parameter (ASCII	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		(ASCII encoded) as defined in GSMP extensions for	ncoded)	

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 that determines the source from which the	integer: 4-octet • 0 = none • 1 = static Layer 2 • 2 = d	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hetransmitspeedisderived.Overridesglobalconfigu		ynamiclayer2.Thismethodisnotsupppor

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ration in the CLI.		ted; the static Layer 2 method is used inst

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			• 3 = CoS. This method is not supported.	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				d;thea ctual method issued inst ead. • 4=

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			actual • 5 = ANCP • 6 = PPPoEIAtags	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-97	IGMP-Immediate-Leave	IGMP Immediate Leave.	integer: 4-octet • 0 = disable • 1 = enable	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.	integer: 4-octet • 0 = disable • 1 = enable	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 IPv6 interfa	string: *policy-name*	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ce.		

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 IPv6 interf	string: *policy-name*	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ace.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-108	CoS-Parameter-Type	CoS traffic-shaping parameter type and description:	Two parts, delimited by white space: • Parameter	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• **T01:** Scheduler-map name • **T02:** Shaping	• Example • **T01smap-type** Parameter value	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• rate • **T03**: Guaranteed rate • **T04**: Delay-	**basic** **T0250m** **T0311m** **T0420000** **T0520000**	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• **T05**: Excess rate • **T06**: Traffic	• buffer rate • **T06**: Traffic	**T06tcp-gold** **T07frame-mode** **T0850**

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			8: Byte adjust • **T09**: Adjustminimum • **T10**	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			:Excess-rate high • **T11**:Excess-rate	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			low • **T12**: Shaping rate burst • **T13**: Guara	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			nteedrateburst	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-109	DHCP-Guided-Relay-Server	IP address of DHCP server that at DHCP relay ya g	integer: 4-byte *ip-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		entusestoforwardthediscoverPDUs.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-110	Acc-Loop-Cir-Id	Identification of the subscriber node connection	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 node.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-112	Acc-Aggr-Cir-Id-Asc	Identification of the uplink on the access node, as	string: up to 63 ASCII characters	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		in the following examples: •	Ethernet access sag	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			gre gation_ethernet slot/port[:inne	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			r-vlan-id][:outer-vlan-id]	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	ATM aggregation_atm slot/port:vpi.v	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			ci	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-113	Act-Data-Rate-Up	Actual upstream data rate of the subscriber's sy	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		nchronizedDSLink.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		synchronized DSL link.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-115	Min-Data-Rate-Up	Minimum upstream data rate configured 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 downstream data rate configured for port	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
26-117	Att-Data-Rate-Up	Maximum upstream data rate that the subscribe	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rcantattain.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-118	Att-Data-Rate-Dn	Maximum downstream data rate that the subscr	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		iberca natt ain.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-119	Max-Data-Rate-Up	Maximum upstream data rate configured 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
26-120	Max-Data-Rate-Dn	Maximum downstream data rate configured for port	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 upstream data rate in low power state co	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-122	Min-LP-Data-Rate-Dn	Minimum downstream data 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
		configured 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 one-way upstream interleaving delay co	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

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 one-way upstream interleaving	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gdelay.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-125	Max-Interlv-Delay-Dn	Maximum one-way downstream interleaving delay	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		configured 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 one-way downstream interleaving	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		vingdelay.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-127	DSL-Line-State	State of the DSL line.	integer: 4-octet • 1 = Showuptime • 2 = Idle • 3 = Sil	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				ent

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-128	DSL-Type	Encapsulation used by the subscriber associate	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dwith the DSLAM interface from which requests are		

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 set to apply to the dynamic profile.	string: *interface-set-name*	No

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 interim accounting updates	• range=600through86400seconds • 0=disable	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tes for this service. Tagged VSA, which supports	**NOTE:** Values are rounded up to the next higher mult	abled

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		8 tags (1-8).	iple of 10 minutes. For example, a setting of 900 seconds (15 min	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			utes) is rounded up to 20 minutes (1200 seconds).	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-141	Downstream-Calculated-QoS-Rate	Calculated (adjusted) downstream QoS rate in Kbps	range = 1000 through 4,294,967,295	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		eresultsinanimmediateInterim-Accountingreque		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	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 (adjusted) upstream QoS rate in Kbps as	range = 1000 through 4,294,967,295	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 er		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		esultsinanimmediateInterim-Accountingrequ e		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	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 allowable client sessions per interface. Fo	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rDHCP clients, this value is the maximum sessions pe		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rlogicalinterface.ForPPPoEclients,thisvalueisthe		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		maximum sessions (PPPoE interfaces) per PPPoE u		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-146	CoS-Scheduler-Pmt-Type	CoS scheduler parameter type and description: Value: Three parts, delimited by white space: • Scheduler		Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• Null: CoS scheduler name • **T01**: CoS sched Ex	name • Parameter type • Parameter value	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	unspecified: • transmit rate • **T02**: CoS scheduler	**be-sched** **be-schedT012m** **be-sc**

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• **T03**: CoS scheduler priority	uffersize **hedT0226**	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	**T04**: CoS scheduler drop-profile low • **T0**	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			5:CoSschedulerdrop-profilemedium-	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	low **T06**:CoSschedulerdrop-profileme	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		• **T07**: CoS scheduler drop-p	dium-high	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		•	rofilehigh **T08**:CoSschedulerdrop-pr	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			of file any	
26-151	IPv6-Acct-Input-Octets	IPv6 receive octets.	integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-152	IPv6-Acct-Output-Octets	IPv6 transmit octets.	integer	No
26-153	IPv6-Acct-Input-Packets	IPv6 receive 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 transmit packets.	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 receive 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 transmit gigawords.	integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-158	PPPoE-Padn	Route add for PPPoE sessions	string	No

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-facin	integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gtrunkphysicalinterface. Vlan-Map-Id (26-160),		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Inner-Vlan-Map-Id (26-184), and Coree-Facing-In		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		workserviceprovider-facinglocationforthesubs		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		cribe 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
		olesaleconfiguration.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-161	IPv6-Delegated-Pool-Name	Address pool used to locally allocate address delegated	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ion.		

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 connect	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ion.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-164	IPv4-Release-Control	Indicates to server status of on-demand addr	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		essallocationanddeallocation.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-173	Service-Activate-Type	Indication of service activation type. This is a tagge	integer: 4-octet • 1 = dynamic-profile for ressi	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dattribute.		dentialservices • 2=op-scriptforbusi

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				nesservices

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-174	Client-Profile-Name	Enables RADIUS to override an assigned client dy	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		namic profile with the included *client-profile-name*-		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		string. Enables RADIUS to distinguish different dyn		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		amic profiles used on the router when the *version*		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		-*alias-string* is included.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-177	Cos-Shaping-Rate	Effectivedownstreamshapingratefor subscrib	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		er.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-178	Action-Reason	Indicates detailed reason for CoA response.	String: "100 In progress" "104 Service active" "120 Service en	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		Typically used for CoA failures.	otfound" "122 Execution failure" "105 Initial processing per	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			formed" "123 No services" "124 Services limit exceeded"	

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			concurrentCoAs" "126CoARequestTimeout" "127Log	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			out in 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 can use these	integer • range=0 through 16777721 54 GB units	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rvice;serviceisdeactivatedwhenthevolumeisexc	• 0=nolimit	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		eeded.TaggedVSA, whichsupports8tags(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 value of service and time quota for re existin	string: *service-name*	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		gservice.TaggedVSA, 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
26-181	DHCPv6-Guided-Relay-Server	IPv6 address of DHCPv6 server to which DH	hexadecimal string: *ipv6-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		CPv6 relay agent forwards the Solicit and subse		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		quentPDUs.Usemultipleinstancesofthe VSAtospe		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		cify a list of servers.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ibute.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-183	Acc-Loop-Encap	Reports the ANCP Access-Loop-Encapsulation at	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 specified provisi	integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		one dont the core-facing physical interface, used t		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		oswap(replace)theautosensedVLANtagonthe		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		access interface. Vlan-Map-Id (26-160), Inner-V		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		lan-Map-Id(26-184),andCore-Facing-Interfac		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		erviceprovider-facinglocationforthesubscriberf		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		orthe Layer 2 cross-connectina Layer 2 wholesale		

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	Name of the core-facing physical interface that	string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		forwards the Layer 2 wholesale session's downst		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		reamandupstreamtrafficrelativetothenetworkse		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rviceprovider(NSP)router. Vlan-Map-Id(26-		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		160), Inner-Vlan-Map-Id (26-184), and Core-Faci		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		enetworkserviceprovider-facinglocationforthe		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		subscriber 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
		r2wholesaleconfiguration.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-189	DHCP-First-Relay-IPv4-Address	IPv4 address of the first relay link of a client/se	integer: 4-byte *ip-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		reverbinding.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-190	DHCP-First-Relay-IPv6-Address	IPv6 address of the first relay link of a client/se	hexadecimal string: *ipv6-address*	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		reverbinding.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ilyanyinterface.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-193	Pim-Enable	Enable or disable PIM on a BRAS user's interface.	integer: 4-octet • 0 = disable • any non zero value	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				e=enable

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-194	Bulk-CoA-Transaction-Id	A common identifier to associate the series	integer: 4-octet	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		softrelatedCoARequeststasatransactiontheatrib		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		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	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
		heBulk-CoA-Transaction-IdVSA.Thisattributeisun		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		tagged and the value 0 is reserved.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-196	IPv4-Input-Service-Set	Name of an IPv4 input service set to be attach	string	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ed.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hed.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-198	IPv4-Input-Service-Filter	Name of an IPv4 input service filter to be attach	string	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ed.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ched.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-200	IPv6-Input-Service-Set	Name of an IPv6 input service set to be attach	string	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ed.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hed.		

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 be attach	string	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ed.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-203	IPv6-Output-Service-Filter	Name 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
		ched.		
26-204	Adv-Pcef-Profile-Name	Name of a PCEF profile to be attached.	string	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	Reason that the client application is reauthentica	integer • 0 = disable • 1 = Initiate reauthentic	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ted.		ationwhenDHCPrenewrequestisreceived

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				id

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-207	DHCPv6-Options	DHCPv6 client and server options exchanged wit	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		the RADIUS server as TLV options. In release see		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		arlierthanJunosOSRelease17.4.1R1,thisVSAisno		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		t supported.DHCPv6 options are included instead 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	DHCPv4 packet header sent to the RADIUS Sse	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rver; used to instantiated dynamic subscriber interfa		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ces.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-209	DHCPv6-Header	DHCPv6 packet header sent to the RADIUS Sse	hexadecimal string	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rver; used to instantiated dynamic subscriber interfa		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ces.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-210	Acct-Request-Reason	Reason for sending an Accounting-Request message	integer: 4-octet • 0x00001 = Acct-Start-Ack;	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		sage.		that is, receipt of an Acct-response for t

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				heAcct-StartmessageOx0002=Period

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				ic/Timedinterval interim 0x00004=IPa

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				active 0x00008=IPinactive 0x00010=IPv6ac

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				tive 0x0020=IPv6inactive 0x0040=Sessi

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				onactive 0x0080=Sessionactive 0x010

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				0 = Linespeedchange 0x02000 = Addressas

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				signmentchange 0x04000=Completionofp

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				rocessingofCoArequest

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 VLAN tag	hexadecimal string: • range=0x6000 through 0x	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			• ffff.0x8100=InnerVLANtagfordesignat	

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				edL2BSAsubscribers

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-212	Routing-Services	Determines whether routing service capa	integer: 4-octet • 0x00000=Disable install atio	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		bility is enabled or disabled.	• 0x0001 = Enable in	no routing services.

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
				stallationoffroutingservices. Anyvalueoth

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
			erthan0or1isrejected.	

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 an interface set to associa	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		teitanditsmemberlinkswithhanaggregatedEth		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ernetmemberlinkfortargetedistribution.		

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 an interface to associate it	integer: 4-octet	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		egatedEthernetmemberlinkfortargetedist		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ribution.When an interface set does not have a weig		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ht, then the interface weight value for the first au		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		thorizedsubscriberinterfaceisusedforthes		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		et.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–216	Hybrid-Access-DSL-Downstream-Speed	Specify a downstream bandwidth 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 access tunnel for a subscriber. Used b		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ythePFEforload-balancingtraffic acrossthe DSLa		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ndLTElegs.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–217	Hybrid-Access-LTE-Downstream-Speed	Specify a downstream bandwidth 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 a subscriber. Use		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		dby the Packet Forwarding Engine for load-bala		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ncing traffic across the DSL and LTE legs.		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–218	Connection-Status-Message	Specifies connection parameters as an encoding	string	Yes

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hat is presented to other remote peer/client (such		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		as a home gateway). This is a logical extension to th		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		formatandsemantics. Theauthdprocessuses		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		only the first instance if it receives multiple instances o		

Table 2: Supported Juniper Networks VSAs *(Continued)*

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

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–219	PON-Access-Type	Type of PON transmission system in use: •	32-bit integer 0—OTHER	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-PON • 5—WDM-PON • 7—UNKNOWN	

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) Averaged downstream data 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 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 Kbps		

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 Kbps		

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 the 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 data rate for the P	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	(G.fast) Expected upstream throughput, which is	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		the net data rate reduced 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 upstream throughput, which his	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		the net data rate reduced 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 th	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		rough 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	(G.fast) Maximum attainable expected downstream	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		through put, 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 data rate (net data rat	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, adjusted down by any through p		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		utcapabilitylimitations,inKbps		

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 data rate (net data rar	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		ate) for the local loop, adjusted down by any through		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		hput capability limitations, 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
		etdatarate) for the local loop, adjusted down by an		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		y through put 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.fast) Maximum attainable downstream data rat	32-bit integer	No

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		e (net data rate) for the local loop, adjusted down by		

Table 2: Supported Juniper Networks VSAs *(Continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
		anythroughputcapabilitylimitations,inKbps		

AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 3 on page 434 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

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
1	User-Name	✓	✓	–	–	–	✓
2	User-Password	✓	–	–	–	–	–
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	–	✓	–	–	–	–
12	Framed-MTU	✓	–	–	–	–	–
18	Reply-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
22	Framed-Route	–	✓	–	–	–	–
24	State	✓	✓	–	✓	–	–
25	Class	–	✓	–	–	✓	–
26-1	Virtual-Router	✓	✓	–	–	✓	–
26-4	Primary-DNS	–	✓	–	–	–	–
26-5	Secondary-DNS	–	✓	–	–	–	–
26-6	Primary-WINS	–	✓	–	–	–	–
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	✓	–	–	–	–	–

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
26-25	Redirect-VR-Name	–	✓	–	–	–	–
26-31	Service-Bundle	–	✓	–	–	–	–
26-33	Tunnel-Maximum-Sessions	–	✓	–	–	–	–
26-34	Framed-IP-Route-Tag	–	✓	–	–	–	–
26-47	Ipv6-Primary-DNS	–	✓	–	–	–	–
26-48	Ipv6-Secondary-DNS	–	✓	–	–	–	–
26-55	DHCP-Options	✓	–	–	–	–	–
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	✓	–	–	–	–	–

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
26-64	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	–	✓	–	–	–	–
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	✓	–	–	–	–	–

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
26-93	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	–	✓	–	–	–	–
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	✓	–	–	–	–	–

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
26-113	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 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
26-123	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	✓	–	–	–	–	–
26-130	QoS-Set-Name	–	✓	–	–	–	–
26-140	Service-Interim-Account-Interval	–	✓	–	–	✓	–
26-141	Downstream-Calculated-QoS-Rate	✓	–	–	–	–	–
26-142	Upstream-Calculated-QoS-Rate	✓	–	–	–	–	–

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
26-143	Max-Clients-Per-Interface	–	✓	–	–	–	–
26-146	Cos-Scheduler-Pmt-Type	–	✓	–	–	✓	–
26-158	PPPoE-Padn	–	✓	–	–	–	–
26-160	Vlan-Map-Id	–	✓	–	–	–	–
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	–	✓	–	–	–	–

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
26-179	Service-Volume-Gigawords	–	✓	–	–	✓	–
26-180	Update-Service	–	–	–	–	✓	–
26-181	DHCPv6-Guided-Relay-Server	–	✓	–	–	–	–
26-182	Acc-Loop-Remote-Id	✓	–	–	–	–	–
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	✓	–	–	–	✓	–

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
26-193	Pim-Enable	–	✓	–	–	–	–
26-194	Bulk-CoA-Transaction-Id	–	–	–	–	✓	–
26-195	Bulk-CoA-Identifier	–	–	–	–	✓	–
26-196	IPv4-Input-Service-Set	✓	–	–	–	–	–
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	✓	–	–	–	–	–

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
26-203	IPv6-Output-Service-Filter	✓	–	–	–	–	–
26-204	Adv-Pcef-Profile-Name	✓	–	–	–	–	–
26-205	Adv-Pcef-Rule-Name	✓	–	–	–	–	–
26-206	Re-Authentication-On-Renew	–	✓	–	–	–	–
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	–	✓	–	–	–	–

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
26-214	Interface-Targeting-Weight	–	✓	–	–	–	–
26–216	Hybrid-Access-DSL-Downstream-Speed	–	✓	–	–	–	–
26-217	Hybrid-Access-LTE-Downstream-Speed	–	✓	–	–	–	–
26–218	Connection-Status-Message	–	✓	–	–	✓	–
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	✓	–	–	–	–	–

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
26–225	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	✓	–	–	–	–	–
26–232	Attainable-Gamma-Data-Rate-Upstream	✓	–	–	–	–	–
26–233	Attainable-Gamma-Data-Rate-Downstream	✓	–	–	–	–	–
27	Session-Timeout	–	✓	–	✓	✓	–
28	Idle-Timeout	–	✓	–	✓	–	–

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
31	Calling-Station-ID	✓	–	–	–	✓	–
32	NAS-Identifier	✓	–	–	–	–	–
44	Acct-Session-ID	✓	–	–	–	✓	✓
61	NAS-Port-Type	✓	–	–	–	–	–
64	Tunnel-Type	✓	✓	–	–	–	–
65	Tunnel-Medium-Type	✓	✓	–	–	–	–
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	✓	–	–	–	✓	–

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
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	–	✓	–	–	–	–
98	Login-IPv6-Host	✓	✓	–	–	–	–
99	Framed-IPv6-Route	–	✓	–	–	–	–
100	Framed-IPv6-Pool	–	✓	–	–	–	–
123	Delegated-IPv6-Prefix	–	✓	–	–	–	–
168	Framed-IP-Address	–	✓	–	–	–	–
242	Ascend-Data-Filter	–	✓	–	–	✓	–

AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 4 on page 449 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	✓	✓	✓	–	–
26-24	PPPoE-Description	✓	✓	✓	–	–

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
26-42	Input-Gigapackets	–	✓	✓	–	–
26-43	Output-Gigapackets	–	✓	✓	–	–
26-47	Ipv6-Primary-DNS	✓	✓	✓	–	–
26-48	Ipv6-Secondary-DNS	✓	✓	✓	–	–
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	–	✓	✓	–	–
26-92	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	✓	✓	✓	–	–

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
26-113	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	✓	✓	✓	–	–
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 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
26-128	DSL-Type	✓	✓	✓	–	–
26-141	Downstream-Calculated-QoS-Rate	✓	✓	✓	–	–
26-142	Upstream-Calculated-QoS-Rate	✓	✓	✓	–	–
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	–	–	✓	–	–
26-177	Cos-Shaping-Rate	✓	✓	✓	–	–

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
26-182	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	✓	–	✓	–	–
26–219	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	✓	✓	✓	–	–

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
26–223	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	✓	✓	✓	–	–
26–232	Attainable-Gamma-Data-Rate-Upstream	✓	✓	✓	–	–
26–233	Attainable-Gamma-Data-Rate-Downstream	✓	✓	✓	–	–
31	Calling-Station-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
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	–	✓	✓	–	–
49	Acct-Terminate-Cause	–	✓	✓	–	–
52	Acct-Input-Gigawords	–	✓	✓	–	–
53	Acct-Output-Gigawords	–	✓	✓	–	–
55	Event-Timestamp	✓	✓	✓	✓	✓
61	NAS-Port-Type	✓	✓	✓	–	–

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
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	✓	✓	✓	–	–
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 | **464**

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 458 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	Identifier for the subscriber agent circuit ID (ACI) that corresponds to the access node interface from which subscriber requests are initiated. 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.	DSL, PON	string

Table 5: DSL Forum VSAs (Vendor ID 3561) *(Continued)*

Type	Name	Description	Access Type	Value
2	Agent-Remote-Id	Unique identifier for the subscriber associated with the access node interface from which requests are initiated. 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.	DSL, PON	string
3	Access-Aggregation-Circuit-ID-ASCII	ASCII identifier for the subscriber access line, based on its network-facing logical appearance If the string begins with a # sign, then the remainder of the string represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. The string is used as the name of a CoS Level 2 interface set that groups subscribers.	DSL, PON	string
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

Table 5: DSL Forum VSAs (Vendor ID 3561) ***(Continued)***

Type	Name	Description	Access Type	Value
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
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

Table 5: DSL Forum VSAs (Vendor ID 3561) ***(Continued)***

Type	Name	Description	Access Type	Value
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
145	DSL-Type	Type of DSL transmission system in use: • 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: • 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 464 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 467 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 *primary-dns-address*
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 *secondary-dns-address*

SEE ALSO

DNS Address Assignment Precedence | 778

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 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 583 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]
- Junos OS Release 18.2 Subscriber Management RADIUS Dictionary [DCT]
- Junos OS Release 17.4 Subscriber Management RADIUS Dictionary [DCT]
- Junos OS Release 17.1 Subscriber Management RADIUS Dictionary [DCT]
- Junos OS Release 16.2 Subscriber Management RADIUS Dictionary [DCT]
- Junos OS Release 16.1 Subscriber Management RADIUS Dictionary [DCT]
- Junos OS Release 15.1 Subscriber Management RADIUS Dictionary [DCT]

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 | **476**

Configuring a Calling-Station-ID with Additional Options | **490**

Release History Table

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 | **550**
- RADIUS NAS Port Attributes and Options | **518**

CHAPTER 2

RADIUS for Subscriber Management

IN THIS CHAPTER

- RADIUS Servers and Parameters for Subscriber Access | **476**
- Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers | **498**
- Session Options for Subscriber Access | **503**
- RADIUS NAS Port Attributes and Options | **518**
- RADIUS Logical Line Identification | **543**
- RADIUS Authentication and Accounting Basic Configuration | **550**
- RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | **556**
- Configuring RADIUS Reauthentication for DHCP Subscribers | **567**
- RADIUS Accounting for Subscriber Access | **570**
- Verifying and Managing Subscriber AAA Information | **601**
- Session Termination Causes and RADIUS Termination Cause Codes | **603**
- AAA Termination Causes and Code Values | **608**
- DHCP Termination Causes and Code Values | **610**
- L2TP Termination Causes and Code Values | **611**
- PPP Termination Causes and Code Values | **638**
- VLAN Termination Causes and Code Values | **651**

RADIUS Servers and Parameters for Subscriber Access

IN THIS SECTION

- RADIUS Authentication and Accounting Server Definition | **477**
- Configuring Options that Apply to All RADIUS Servers | **480**

- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | **482**
- Configuring Access Profile Options for Interactions with RADIUS Servers | **483**
- Configuring a Calling-Station-ID with Additional Options | **490**
- Filtering RADIUS Attributes and VSAs from RADIUS Messages | **494**

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 543.
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" on page 2483 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 483.

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 482 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 523 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" on page 476.

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 490 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:

a. Delete the `juniper-access-line-attributes` option from all access profiles that include it.

b. Perform the software downgrade.

c. 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 520 and "Configuring the Order in Which Optional Values Appear in the NAS-Port-ID" on page 521 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 that Apply to All 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 ]
  ```

Release History Table

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 | 498
- Interface Description Storage and Reporting Configuration | 503

Interface Description Storage and Reporting Overview

IN THIS SECTION

- Interface Description Precedence | 498
- Example: Reporting Interface Descriptions on Non-Underlying Logical Interfaces | 499
- Reporting Interface Descriptions on Underlying Logical Interfaces | 500
- Example: PPP over an Underlying VLAN Demux Interface | 500
- Example: Reporting Interface Descriptions on Dynamic VLANs | 502

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 | 476

Session Options for Subscriber Access

IN THIS SECTION

- Understanding Session Options for Subscriber Access | 504
- Subscriber Session Timeout Options | 511
- Limiting the Number of Active Sessions per Username and Access Profile | 512
- Configuring Username Modification for Subscriber Sessions | 513
- Removing Inactive Dynamic Subscriber VLANs | 516

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 | 504
- Limits on Subscriber Sessions per Username and Access Profile | 507
- Benefits of Limiting Sessions for Usernames with the CLI | 509
- Subscriber Username Modification | 509
- Benefits of Subscriber Username Modification | 510

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@example.com	user1
right-to-left	user1@example.com	user1

g043376

- 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

- 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

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.

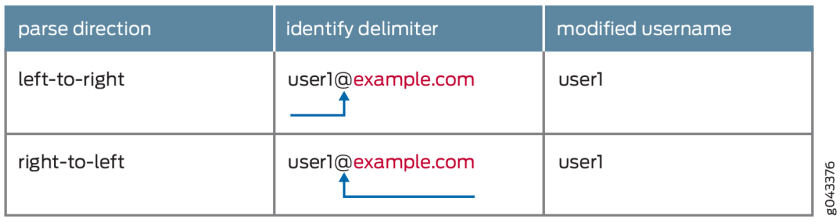

Now suppose the subscriber logs in with the username, user1@test@example.com. For a username like this, the parsing direction makes a difference in the modified username. The configuration determines that the first instance of the delimiter @ is found first, because the name is parsed from left to right. 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@test@example.com	user1
right-to-left	user1@test@example.com	user1@test

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

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
```

Release History Table

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 | 476

RADIUS NAS Port Attributes and Options

IN THIS SECTION

- Manual Configuration of the NAS-Port-ID RADIUS Attribute | **518**
- Configuring a NAS-Port-ID with Additional Options | **520**
- Configuring the Order in Which Optional Values Appear in the NAS-Port-ID | **521**
- Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers | **523**
- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview | **524**
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **526**
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **527**
- Manual Configuration of the NAS-Port-Type RADIUS Attribute | **528**
- Configuring the RADIUS NAS-Port-Type per Physical Interface | **531**
- Configuring the RADIUS NAS-Port-Type per VLAN | **532**
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN | **534**
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface | **536**
- Configuring the RADIUS NAS-Port Extended Format per VLAN | **537**
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN | **539**
- Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces | **541**

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 521.

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/adapter/port/SVLAN ID/VLAN ID`
- Subscribers over ATM interfaces—combination of `slot/adapter/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 | 525
- NAS-Port RADIUS Attribute | 525

NAS-Port Options Configuration and Subscriber Network Access Models | 525
NAS-Port Options Definition | 526

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 531.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per VLAN" on page 532.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per Stacked VLAN" on page 534.
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 536.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per VLAN" on page 537.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN" on page 539.

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 529 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
value	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;
    }
}
```

Release History Table

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 | 483
- RADIUS Servers and Parameters for Subscriber Access | 476

RADIUS Logical Line Identification

IN THIS SECTION

- RADIUS Logical Line Identifier (LLID) Overview | 544
- RADIUS Attributes for LLID Preauthentication Requests | 545
- Configuring Logical Line Identification (LLID) Preauthentication | 546
- Configuring a Port and Password for LLID Preauthentication Requests | 548
- Verifying and Managing LLID Preauthentication Configuration | 549

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 545 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	(Non-standard use of attribute.) Identifying information for the user associated with the LLID, in the following format. `nas-port:`*`nas-ip-address`*`:`*`nas-port-id`* Example: `nas-port:198.51.100.117:ge-1/0/5:100` **NOTE**: The router strips any dynamically generated information from the User-Name attribute during preauthentication.

Table 10: RADIUS Attributes for LLID Preauthentication Requests *(Continued)*

Attribute Number	Attribute Name	Description
2	User-Password	(Non-standard use of attribute.) Password of the user to be authenticated. Example: Always set to `juniper`
4	NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user Example: `198.51.100.117`
5	NAS-Port	Physical port number of the NAS that is authenticating the user. Always interpreted as a bit field
6	Service-Type	Type of service the user requested or the type of service to be provided. Example: `gold-service`
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	(Non-standard use of attribute.) The user name. Example: `jdoe@xyzcorp.example.com`
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
```

4. 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.

5. (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
```

6. (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 | 550
- Action | 550

Purpose

Display statistics and configuration information related to logical line identification (LLID) preauthenticaion.

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 | **550**
- Specifying the Authentication and Accounting Methods for Subscriber Access | **551**
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access | **552**
- Configuring Local Authentication and Authorization for Subscribers | **552**

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 551.

2. Specify how accounting statistics are collected.

   See "Configuring Per-Subscriber Session Accounting" on page 573.

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.

Release History Table

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 | 476

Configuring Local Authentication and Authorization for Subscribers | 552

Configuring Per-Subscriber Session Accounting | 573

Specifying the Authentication and Accounting Methods for Subscriber Access | 551

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers

IN THIS SECTION

- Benefits of Reauthentication | 557
- Functionality | 557
- Dual-Stack Subscribers | 560
- Packet Flow | 562
- RADIUS Attributes Supported for Reauthentication | 565

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 563, 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 564. 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 558 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.

Table 11: Processing Multiple Request Types *(Continued)*

Request in Progress	Additional Request Received for Same Subscriber	Action
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 serer 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 turns 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 563 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

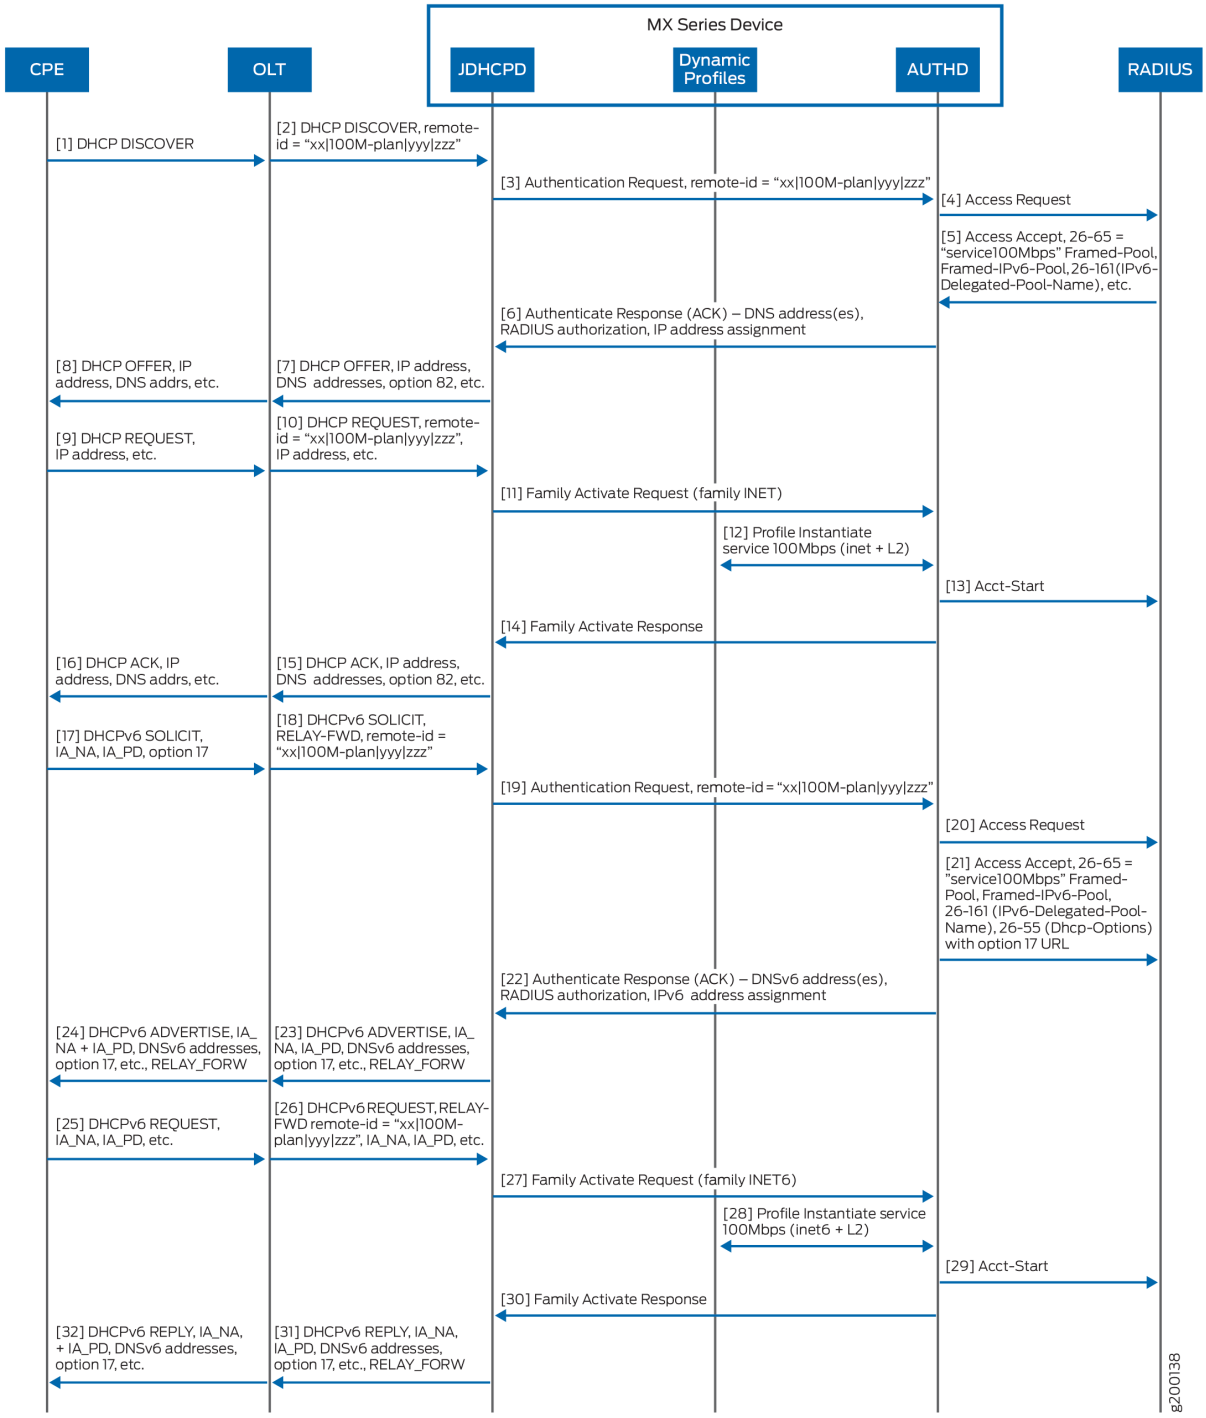

Service Plan Change

Figure 2 on page 564 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

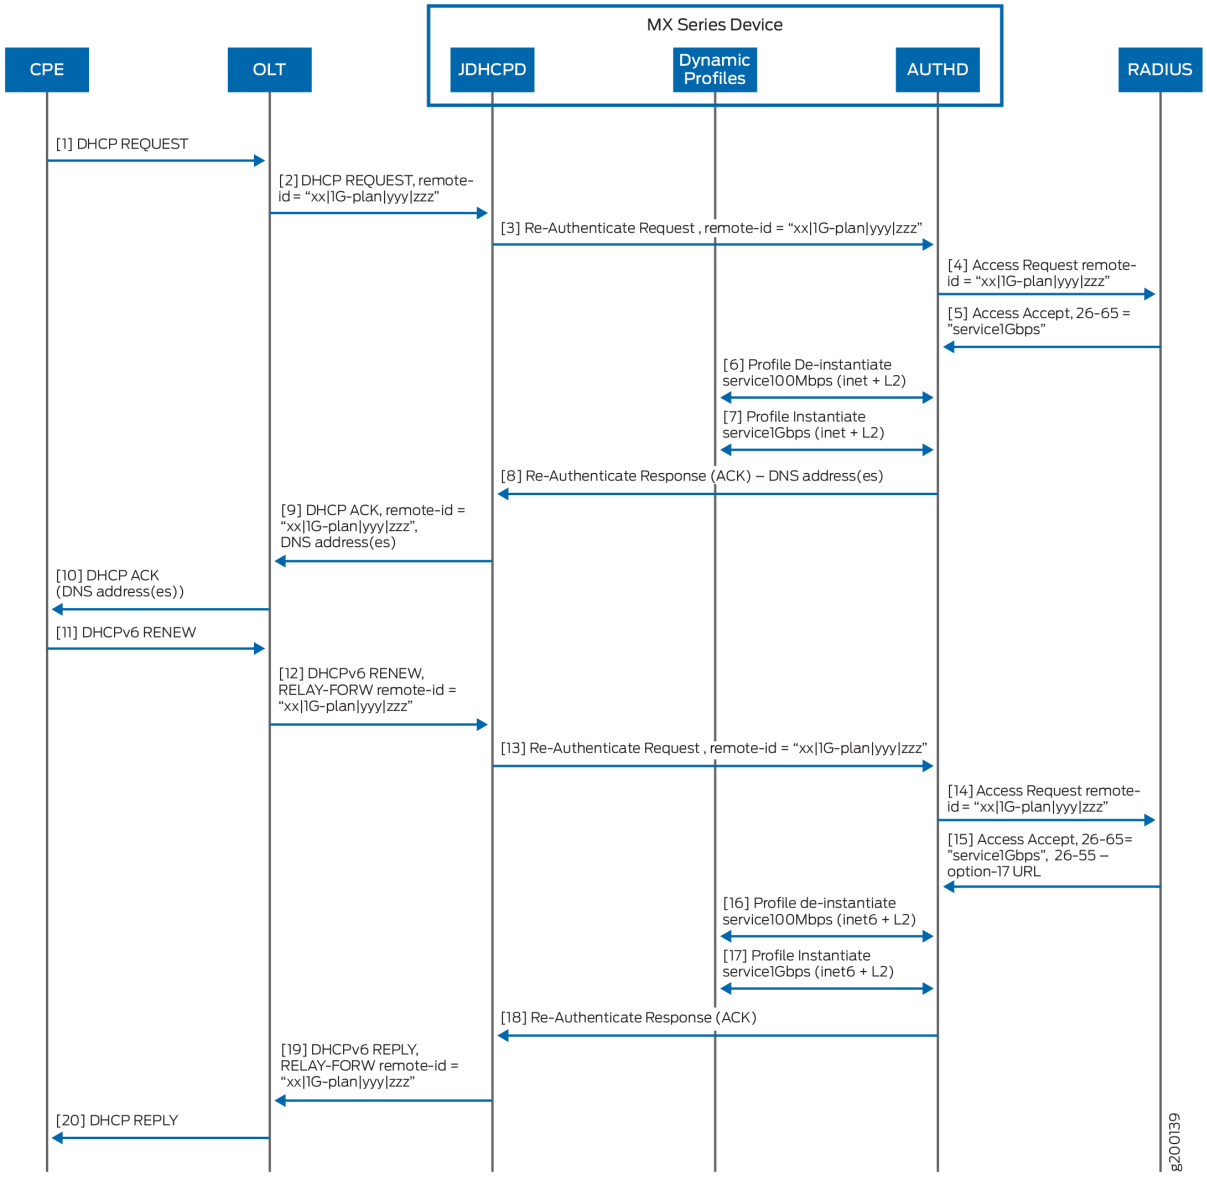

RADIUS Attributes Supported for Reauthentication

Table 12 on page 565 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	The authd process compares the list of services in the VSA to the services that are already active for that subscriber session. • 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. 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.
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	If the value is 1 (enable), authd adds the value to the subscriber session database if it is not already present. If the value is 0 (disable) and a value of 1 is already present in the database, authd sets the database value to 0. 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.
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.

Release History Table

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 | **567**

Single-Session DHCP Dual-Stack Overview | **1004**

Understanding Differences Between Legacy DHCP and Extended DHCP | **691**

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.

Release History Table

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 | 556

Single-Session DHCP Dual-Stack Overview | 1004

DHCP-Initiated Service Change Based on Remote ID | 744

Understanding Differences Between Legacy DHCP and Extended DHCP | 691

RADIUS Accounting for Subscriber Access

IN THIS SECTION

- RADIUS Accounting Statistics for Subscriber Access Overview | 571
- RADIUS Acct-On and Acct-Off Messages | 572
- Configuring Per-Subscriber Session Accounting | 573
- Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI | 576
- Understanding RADIUS Accounting Duplicate Reporting | 578
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting | 580
- Configuring Per-Service Session Accounting | 581
- Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 583
- Configuring Service Packet Counting for Volume Statistics | 585
- Configuring Service Accounting | 586
- Preservation of RADIUS Accounting Information During an Accounting Server Outage | 588
- Configuring Back-up Options for RADIUS Accounting | 591
- Forcing the Router to Contact the Accounting Server Immediately | 592
- Monitoring Pending RADIUS Accounting Stop Messages | 593
- Suspending RADIUS Accounting and Baselining Accounting Statistics Overview | 595
- Configuring RADIUS Accounting Suspension and Baselining Accounting Statistics | 599

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 550.

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 571 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

Table 13: RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting *(Continued)*

Attribute Number	Attribute Name	Type of Statistics
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 | 550

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 in 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 for Junos OS | 448

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 571 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 | 1325

RADIUS Authentication and Accounting Basic Configuration | 550

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 | 578
- Other Scenarios | 579
- Filters for Duplicate Accounting Reports | 580

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 579 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" on page 1872 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 571 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 582 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 585.

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	• 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	• range = 600–86400 seconds • 0 = disabled **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).

SEE ALSO

RADIUS Authentication and Accounting Basic Configuration | 550

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 top 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 uses RADIUS CoAs for service provisioning but does not support Juniper Networks VSAs. For more information about an this use case, see "Processing Cisco VSAs in RADIUS Messages for Service Provisioning" on page 583.

Table 16 on page 587 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 587 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 588 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 `statistics` = `time`	Acct-Interim-Interval (85) = 600 Service-Statistics (26-69) not set	600 time
`update-interval` = 400 `statistics` = `time`	Acct-Interim-Interval (85) not set Service-Statistics (26-69) = 2, time and volume	400 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 589, 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

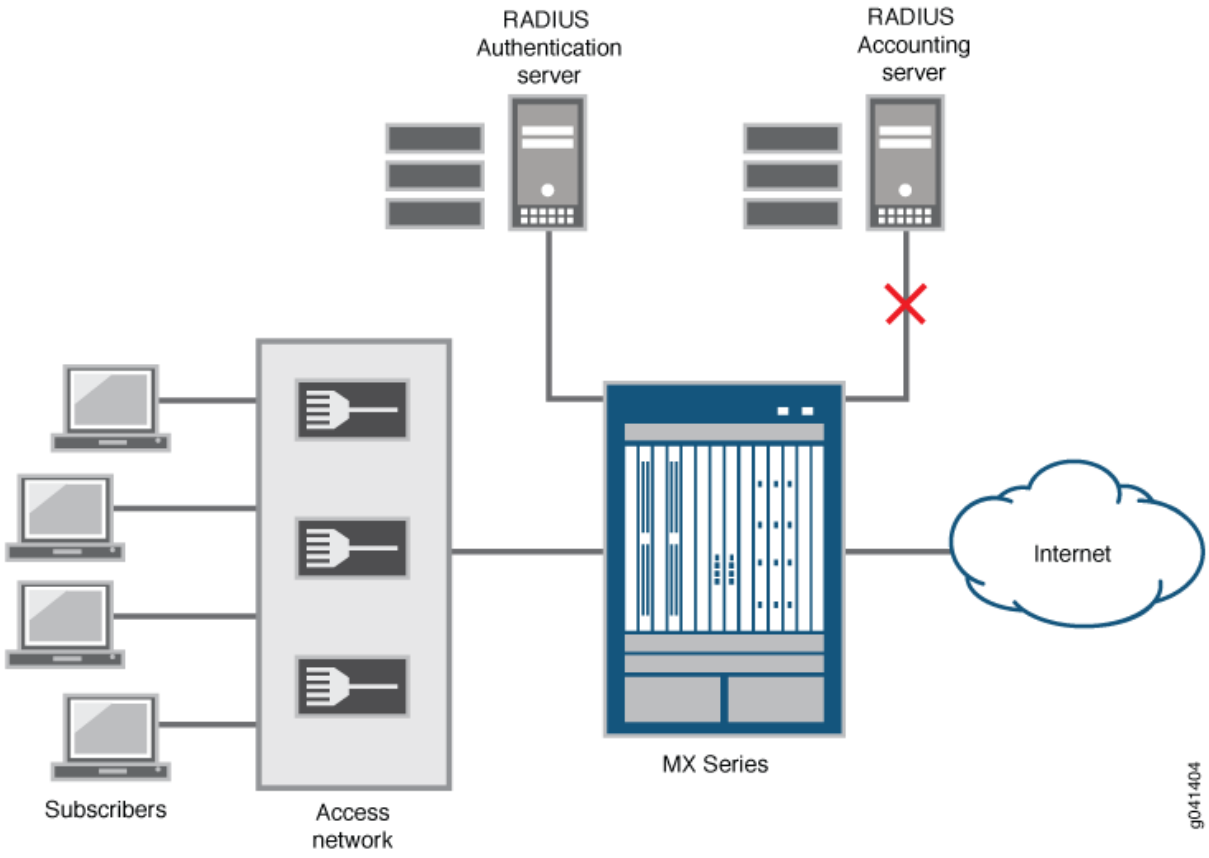

By default, the router must wait until the revert timer expires before it can attempt to contact the non-responsive server again. However, when you configure accounting backup, the revert timer is disabled and the router immediately retries its accounting requests as soon as the router fails to receive 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 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" on page 2958 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" on page 2958 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 | 593
- Action | 593

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 Baselining Accounting Statistics Overview

IN THIS SECTION

- Sequence of Events During the Suspension, Baselining, and Resumption of Accounting | **597**
- Guidelines for Accounting Suspension and Baselining of Statistics | **598**
- Sample Scenarios of Subscriber Accounting Suspension and Baselining | **598**

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, Baselining, 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 Baselining 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 baselining is not supported when accounting is not suspended.
- You cannot specify more than one baseline request during an accounting suspension.
- Baselining 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 baselining 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 baselining 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 baselining is in progress, the state of suspension or baselining 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 Baselining

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. Baselining 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. Baselining 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 599 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.

Release History Table

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 | 550

Standard and Vendor-Specific RADIUS Attributes | 3

Verifying and Managing Subscriber AAA Information

IN THIS SECTION

Purpose | 601

Action | 601

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 | 476

Session Termination Causes and RADIUS Termination Cause Codes

IN THIS SECTION

- Understanding Session Termination Causes and RADIUS Termination Cause Codes | 603
- Mapping Session Termination Causes to Custom Termination Cause Codes | 606

Understanding Session Termination Causes and RADIUS Termination Cause Codes

IN THIS SECTION

- Benefits of Session and Service Termination Cause Codes | 606

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 604 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).
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.

Table 20: RFC-Defined Code Values and Termination Causes *(Continued)*

Code Value	Termination Cause	Description
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.
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.

Table 20: RFC-Defined Code Values and Termination Causes *(Continued)*

Code Value	Termination Cause	Description
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 | 608

DHCP Termination Causes and Code Values | 610

L2TP Termination Causes and Code Values | 611

PPP Termination Causes and Code Values | 638

VLAN Termination Causes and Code Values | 651

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 608 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.
deny-no-resources	10	Subscriber access denied for reasons such as no RADIUS server exists.

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-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 dhcp clear binding command [I do not know the exact syntax])
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.

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-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 | 603

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 611 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 | 603

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 612 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
session call failed	15	Service Unavailable
session create failed limit reached	9	NAS Error
session create failed no resources	9	NAS 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 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
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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
session rx ocrp avp unknown	10	NAS Request
session rx ocrp no resources	10	NAS Request
session rx ocrp unexpected	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 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
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

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 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
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

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 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
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

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 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
session unknown cause	9	NAS Error
session upper create failed	9	NAS Error
session upper removed	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
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
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

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 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 scccn authenticate failed challenge	17	User Error
tunnel rx scccn avp bad hidden	15	Service Unavailable
tunnel rx scccn avp bad value challenge response	15	Service Unavailable
tunnel rx scccn avp malformed bad length	15	Service Unavailable
tunnel rx scccn avp malformed truncated	15	Service Unavailable
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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 scccn authenticate failed challenge	17	User Error
tunnel rx recovery scccn avp bad hidden	15	Service Unavailable
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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
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

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 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
tunnel timeout connection	15	Service Unavailable
tunnel timeout connection recovery tunnel	15	Service Unavailable
tunnel timeout idle	1	User 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
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 | 603

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 639 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 639 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
authenticate pap peer authenticator timeout	17	User 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
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 mped not set yet	10	NAS Request
bundle fail local mrru mismatch	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 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
ip link down	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 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
ip no service	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 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
ipv6 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
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
lcp configured mrru too small	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 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 localaccm	10	NAS Request
lcp no localacfc	10	NAS Request
lcp no local authentication	10	NAS Request
lcp no local endpoint discriminator	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 magic number	10	NAS Request
lcp no local mrru	10	NAS Request
lcp no local mru	10	NAS Request
lcp no localpfc	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
lcp peer terminate term req	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 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
mpls no service	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 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
no ncps available	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
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
osi peer terminate term ack	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 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 | 603

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 651 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 | **603**

CHAPTER 3

Domain Maps for Subscriber Management

IN THIS CHAPTER

- Mapping Subscriber Domains to Access and Session Options | **654**
- Verifying Domain Maps | **673**

Mapping Subscriber Domains to Access and Session Options

IN THIS SECTION

- Domain Mapping Overview | **655**
- Configuring a Domain Map | **659**
- Configuring a Wildcard Domain Map | **661**
- Specifying an Access Profile in a Domain Map | **662**
- Specifying an Address Pool in a Domain Map | **663**
- Specifying a Dynamic Profile in a Domain Map | **664**
- Specifying an AAA Logical System/Routing Instance in a Domain Map | **664**
- Specifying a Target Logical System/Routing Instance in a Domain Map | **665**
- Specifying a Tunnel Profile in a Domain Map | **666**
- Specifying a Tunnel Switch Profile in a Domain Map | **667**
- Configuring Domain and Realm Name Usage for Domain Maps | **667**
- Specifying Domain and Realm Name Delimiters | **668**
- Specifying the Parsing Order for Domain and Realm Names | **669**
- Specifying the Parsing Direction for Domain and Realm Names | **670**
- Enabling Domain Name Stripping | **671**
- Changing the Username and Password to Simplify Off-Chassis Provisioning | **671**

Domain Mapping Overview

IN THIS SECTION

- Types of Domain Maps and Their Order of Precedence | **657**
- Wildcard Domain Map | **657**
- Default Domain Map | **657**
- Domain Map for Subscriber Usernames With No Domain or Realm Name | **658**
- Understanding Domain Maps and Logical System/Routing Instance Contexts | **658**
- Benefits of Using Domain Maps | **659**

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 656 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

- Domains 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 661.

- 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 662.

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 664.

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 663.

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 664.

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 665.

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 667.

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 671.

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 671.

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 671.

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 666.

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 667.

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 659.

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 662.

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 663.

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 664.

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 669.

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 668.

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 670.

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 671.

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 the 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
```

Release History Table

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 | 673

Verifying Domain Maps

IN THIS SECTION

Purpose | 674
Action | 674

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 | 655

Configuring a Domain Map | 659

CHAPTER 4

Testing and Troubleshooting AAA

IN THIS CHAPTER

- AAA Testing and Troubleshooting | 675
- Tracing General Authentication Service (authd) Events for Troubleshooting | 683

AAA Testing and Troubleshooting

IN THIS SECTION

- AAA Configuration Testing and Troubleshooting | 675
- Testing a Subscriber AAA Configuration | 676

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 | 676
- Action | 676

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>
```

Release History Table

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 | 684
- Configuring the Number and Size of General Authentication Service Log Files | 684
- Configuring Access to the General Authentication Service Log File | 685
- Configuring a Regular Expression for General Authentication Service Messages to Be Logged | 685
- Configuring Subscriber Filtering for General Authentication Service Tracing | 686
- Configuring the General Authentication Service Tracing Flags | 687

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
```

Release History Table

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

2

PART

DHCP and DHCPv6 for Subscriber Management

DHCP for Subscriber Management | 690

DHCPv6 for Subscriber Management | 910

CHAPTER 5

DHCP for Subscriber Management

IN THIS CHAPTER

- DHCP Overview | **691**
- DHCP Access Profiles for Subscriber Authentication and Accounting Parameters | **702**
- Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | **706**
- Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | **719**
- DHCP Options and Selective Traffic Processing | **724**
- Using DHCP Option 82 Information | **751**
- Default Services for DHCP Subscribers | **764**
- DHCP Client Attribute and Address Assignment | **766**
- DHCP Lease Times for IP Addresses | **780**
- DHCP Leasequery Methods | **788**
- DHCP Client Authentication With An External AAA Authentication Service | **833**
- Receiving DHCP Options From a RADIUS Server | **838**
- Common DHCP Configuration for Interface Groups and Server Groups | **852**
- Number of DHCP Clients Per Interface | **862**
- Maintaining DHCP Subscribers During Interface Delete Events | **865**
- Dynamic Reconfiguration of Clients From a DHCP Local Server | **870**
- Conserving IP Addresses Using DHCP Auto Logout | **878**
- DHCP Short Cycle Protection | **885**
- DHCP Monitoring and Management | **895**

DHCP Overview

IN THIS SECTION

- Understanding Differences Between Legacy DHCP and Extended DHCP | **691**
- Extended DHCP Relay Agent Overview | **695**
- DHCP Relay Proxy Overview | **698**
- Minimum DHCP Relay Agent Configuration | **700**
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers | **700**

Understanding Differences Between Legacy DHCP and Extended DHCP

IN THIS SECTION

- New Features and Enhancements in Extended DHCP | **691**
- Benefits of Extended DHCP | **693**
- Change in Configuring DHCP Local Server in Extended DHCP Environment | **693**
- Legacy DHCP and Extended DHCP Server Hierarchy Levels Changes | **694**

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. aAny 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 EX Series Switch Software Features Overview for a list of switches that support extended DHCP (VR-aware DHCP).

Table 30 on page 693 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 694:

Table 31: Legacy DHCP and Extended DHCP Server Hierarchy Levels

DHCP Service	Hierarchy
Legacy DHCP server	`edit system services dhcp`
Extended DHCP server	`edit system services dhcp-local-server`
Legacy DHCP relay	`edit forwarding-options helpers bootp`
Extended DHCP relay	`edit forwarding-options dhcp-relay`
Legacy DHCP address pool	`edit system services dhcp pool`
Extended DHCP address pool	`edit access address-assignment pool`

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

Once you enable DHCP-Relay on the MX routers, or QFX or EX switches, the DHCP Snooping feature gets enabled and all DHCP packets incoming through any interface (both configured and unconfigured interface) of the device are analyzed. 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 | **696**
- DHCP Liveness Detection | **697**

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.

NOTE: The PTX Series Packet Transport Routers do not support authentication for DHCP relay agents.

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 916 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 difference 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 | **698**
- Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers | **699**

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.

- 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

This example shows the minimum configuration you need to use the extended DHCP relay agent on the router or switch:

```
[edit forwarding-options]
dhcp-relay {
    server-group {
        test 203.0.113.21;
    }
    active-server-group test;
    group all {
        interface fe-0/0/2.0;
    }
}
```

NOTE: The interface type in this topic is just an example. The `fe-` interface type is not supported by EX Series switches.

The `forward-only` DHCP Relay configurations do not require the S-SA-FP license to be installed.

This example creates a server group and an active server group named `test` with IP address 203.0.113.21. The DHCP relay agent configuration is applied to a group named `all`. Within this group, the DHCP relay agent is enabled on interface fe-0/0/2.0.

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 | 1138

DHCP Client Attribute and Address Assignment | 766

DHCP Client Authentication With An External AAA Authentication Service | 833

DHCP Monitoring and Management | 895

DHCPv6 Local Server | 910

DHCPv6 Relay Agent | 916

DHCPv6 Monitoring and Management | 925

DHCP Access Profiles for Subscriber Authentication and Accounting Parameters

IN THIS SECTION

Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview | 702

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | 703

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 | **703**
- Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients | **704**
- Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces | **705**

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 852.

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-nameaccess-profile profile-name
```

Release History Table

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 | **691**
- DHCPv6 Local Server | **910**
- DHCPv6 Relay Agent | **916**

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings

IN THIS SECTION

- Overriding the Default DHCP Local Server Configuration Settings | **707**
- Overriding the Default DHCP Relay Configuration Settings | **709**
- DHCP Behavior When Renegotiating While in Bound State | **712**
- Sending Release Messages When Clients Are Deleted | **713**
- Disabling Automatic Binding of Stray DHCP Requests | **714**
- Enabling DHCP Relay Proxy Mode | **716**
- Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent | **716**
- Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address | **717**
- Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall | **717**
- Using Layer 2 Unicast Transmission instead of Broadcast for DHCP Packets | **718**
- Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally | **719**

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 862.

- (Optional) Configure DHCP client auto logout.

  See "Automatically Logging Out DHCP Clients" on page 881.

- (Optional) Enable processing of information requests from clients.

  See "Enabling Processing of Client Information Requests" on page 777.

- (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 740.

- (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 955.

- (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.

  See "Configuring DHCPv6 Rapid Commit (MX Series, EX Series)" on page 912.

- (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 1169.

- (Optional, DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

  See "Automatically Logging Out DHCPv6 Clients" on page 884.

- (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 712.

- (Optional, DHCPv4 and DHCPv6) Specify that a short lease be sent to the client.

  See "Configuring DHCP Asymmetric Leasing" on page 786.

- (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 771.

- 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 719.

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 716.

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 716.

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 717.

5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.

   See "Overriding Option 82 Information" on page 751.

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 718.

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 761.

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 862.

9. (DHCPv4 only) Configure client auto logout.

   See "DHCP Auto Logout Overview" on page 879.

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" on page 1704.

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 713.

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 712.

14. (DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

    See "Automatically Logging Out DHCPv6 Clients" on page 884.

15. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.

    See "Disabling DHCP Relay Agent for Interfaces, for Groups, or Globally" on page 719.

16. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.

    See "Disabling Automatic Binding of Stray DHCP Requests" on page 714.

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 1008.

18. (Optional, DHCPv4 and DHCPv6l) Specify that a short lease be sent to the client.

    See "Configuring DHCP Asymmetric Leasing" on page 786.

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 | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916
- Common DHCP Configuration for Interface Groups and Server Groups | 852
- DHCP Options and Selective Traffic Processing | 724
- Using DHCP Option 82 Information | 751

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers

IN THIS SECTION

- Load Balancing DHCP Local Servers by Delaying Responses to Clients | 720
- Configuring a Delay in Local Server Response to DHCP Discover and Solicit Messages | 721

Load Balancing DHCP Local Servers by Delaying Responses to Clients

IN THIS SECTION

- Benefits to Delaying DHCP Local Server Response | 721

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 720 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

Table 32: Characteristics to Identify Clients for Delayed Responses *(Continued)*

Client Characteristic	DHCPv4	DHCPv6
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
```

Release History Table

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 | **706**
- Dynamic Reconfiguration of Clients From a DHCP Local Server | **870**
- DHCP Overview | **691**
- DHCPv6 Local Server | **910**

DHCP Options and Selective Traffic Processing

IN THIS SECTION

- DHCP Options and Selective Traffic Processing Overview | **725**
- Using DHCP Option Information to Selectively Process DHCP Client Traffic | **727**
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings | **728**
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings | **728**
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing | **734**
- DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | **739**
- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | **740**
- DHCP-Initiated Service Change Based on Remote ID | **744**
- Configuring DHCP-Initiated Service Change Based on Remote ID | **745**
- DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | **747**

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 | 729
- Overview | 729
- Configuration | 729
- Verification | 732

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 695.

- (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 852.

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 | **730**
- Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings | **730**

- Results | 731

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 fffff 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 fffff {
                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 | 732

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 | **734**
- Overview | **735**
- Configuration | **735**
- Verification | **738**

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 695 and "DHCPv6 Relay Agent Overview" on page 916.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 852.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 852.

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 | **735**
- Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings | **736**
- Results | **737**

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 | **738**

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 | **742**
- Server-Side Support | **742**
- DHCP Local Server Support | **743**

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 | **745**

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 | **748**
- Processing Order and Actions | **749**
- Benefits of DHCP Relay Forward-Only Action | **750**

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 750. 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

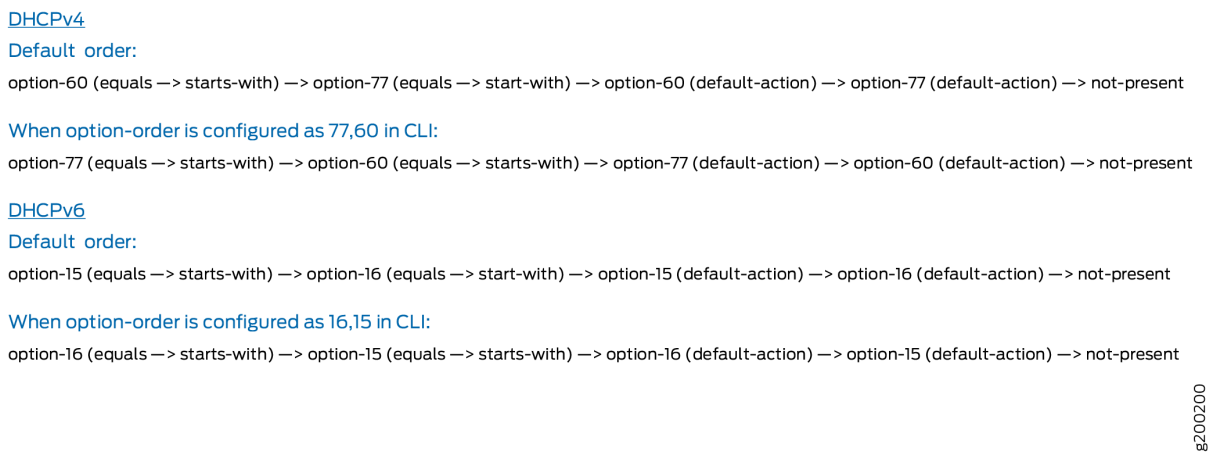

Benefits of DHCP Relay Forward-Only Action

- 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.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, the selective traffic processing feature lets you manage multivendor networks with the extended DHCP and DHCPv6 relay agent.
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 | **691**
DHCPv6 Local Server | **910**
DHCPv6 Relay Agent | **916**

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706
Common DHCP Configuration for Interface Groups and Server Groups | 852
RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 556

Using DHCP Option 82 Information

IN THIS SECTION

- Using DHCP Relay Agent Option 82 Information | 751
- Enable Processing of Untrusted Packets So Option 82 Information Can Be Used | 761
- Extracting an Option 82 or Option 37 Substring to Create an Interface Set | 762

Using DHCP Relay Agent Option 82 Information

IN THIS SECTION

- Configuring Option 82 Information | 752
- Overriding Option 82 Information | 755
- Including a Prefix in DHCP Options | 755
- Including a Textual Description in DHCP Options | 759

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 917.

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 751.

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 751.

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: otl01.xyz101-202
...
Agent Remote ID: ABCD01234|100M|AAAA01234|otl01.xyz101-202
...
```

SEE ALSO

- *Junos OS Predefined Variables*
- *Configuring Predefined Dynamic Variables in Dynamic Profiles*
- *Configuring a Basic Dynamic Profile*

Release History Table

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 | 691

DHCPv6 Local Server | 910

DHCPv6 Relay Agent | 916

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706

Default Services for DHCP Subscribers

IN THIS SECTION

Default Subscriber Service Overview | 764

Configuring a Default Subscriber Service | 765

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 dynamic-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 | **1392**

DHCP Client Attribute and Address Assignment

IN THIS SECTION

- DHCP Attributes Overview | **766**
- Attributes That Can Be Applied to DHCP Clients | **768**
- Configuring DHCP Attributes for All Clients or a Group of Clients | **771**
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | **772**
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **773**
- Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option | **775**
- Specifying the Subnet for DHCP Client Address Assignment | **776**
- DHCP Local Server Handling of Client Information Request Messages | **776**
- Enabling Processing of Client Information Requests | **777**
- DNS Address Assignment Precedence | **778**
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching | **779**

DHCP Attributes Overview

IN THIS SECTION

- Benefits of Configuring DHCP Attributes | **768**

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 768 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
`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

Table 33: DHCP Attributes *(Continued)*

Attribute	Description	DHCP Option
`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
`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 770 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.	–
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.	–

Table 34: DHCPv6 Attributes *(Continued)*

Attribute	Description	DHCPv6 Option
`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 772 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 768 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 1165.

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 do 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 772 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 | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916
- Address-Assignment Pools for Subscriber Management | 1138
- Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706
- Standard and Vendor-Specific RADIUS Attributes | 3

DHCP Lease Times for IP Addresses

IN THIS SECTION

- DHCP Lease Timers | 780
- DHCP Lease-Time Validation Overview | 781
- Configuring a DHCP Lease-Time Threshold | 783
- DHCP Asymmetric Leasing Overview | 784
- Configuring DHCP Asymmetric Leasing | 786

DHCP Lease Timers

IN THIS SECTION

- Benefits of Configuring DHCP Timers in Address Pools | 781

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 | 783

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 782.

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 `ltv-syslog-interval` command	warning	–

Table 35: Lease-Time Violation Event Logging *(Continued)*

Event	syslog	Extended DHCP Traceoptions
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 | **785**

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 709 and "Overriding the Default DHCP Local Server Configuration Settings" on page 707.

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
```

Release History Table

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, *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.
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 | **766**
- Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | **706**
- DHCP Overview | **691**
- DHCPv6 Local Server | **910**
- DHCPv6 Relay Agent | **916**

DHCP Leasequery Methods

IN THIS SECTION

- Benefits of DHCP Leasequery | **789**
- DHCP Individual Leasequery | **790**
- DHCP Bulk Leasequery | **794**
- DHCP Active Leasequery | **799**
- Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations | **811**
- Configuring and Using DHCP Individual Leasequery | **812**
- Configuring and Using DHCP Bulk Leasequery | **814**
- Configuring and Using DHCP Active Leasequery | **818**
- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | **824**
- Verifying and Managing DHCP Individual and Bulk Leasequery Configurations | **829**
- Verifying and Managing DHCP Active Leasequery Operations | **830**

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 | **790**
- DHCPv6 Individual Leasequery | **792**

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 serer 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 791 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 793 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 793 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 | **794**
- DHCPv6 Bulk Leasequery | **797**

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 a 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 788. 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 796 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.
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 796 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 798 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	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. 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. When the query fails, a single LEASEQUERY-REPLY is returned with no binding information.
LEASEQUERY-DONE	16	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. The TCP connection between the relay agent and server is closed when this message is received.
LEASEQUERY-DATA	17	Response from the local server with information about the leases for a single DHCPv6 client or about prefix delegation bindings on a single link. 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.

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 788. The messages can also include the status code added for bulk leasequery described in Table 42 on page 799.

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 | **800**
- DHCPv6 Active Leasequery | **802**
- Chassis-Level Redundancy with Active Leasequery | **804**
- Interface-Level Redundancy with Active Leasequery Topology Discovery | **805**

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 788. 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 801 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	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. Can also be sent between peer relay agents to provide hot standby redundancy for binding information.
DHCPLEASEQUERYSTATUS	17	Response from the local server when it has returned binding information associated with the request. 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.

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 788. The messages can also include the status codes added for active leasequery described inTable 44 on page 802.

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.
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 802 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	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. Can also be sent between peer relay agents to provide hot standby redundancy for binding information.

Table 45: DHCPv6 Active Leasequery Message Types *(Continued)*

Message Type	DHCPv6 Type Value	Description
LEASEQUERY-REPLY	15	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. 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. When the query fails, a single LEASEQUERY-REPLY is returned with no binding information.
LEASEQUERY-DONE	16	Response from the local server that indicates that the connection should be terminated. For example, the sever can send this with a QueryTerminated status code (11) when the server is being shut down.
LEASEQUERY-DATA	17	Response from the local server with information about the leases for a single DHCPv6 client or about prefix delegation bindings on a single link. 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.

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 788 and "DHCPv6 Bulk Leasequery Status Code" on page 788, respectively. The messages can also include the status codes added for active leasequery described in Table 46 on page 803.

Table 46: DHCPv6 Active Leasequery Status Codes

Code	Status	Description
12	DataMissing	The requested binding information is not available.

Table 46: DHCPv6 Active Leasequery Status Codes *(Continued)*

Code	Status	Description
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 804 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 801 and Table 45 on page 802. 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

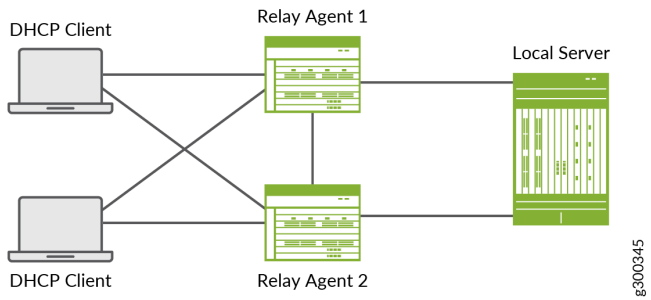

The following sequence describes how the relay agents establish the peer relationship and share binding information when active leasequery is configured on both. This example is for DHCPv4, but the mechanism is the same for DHCPv6.

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 1173 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 807.

Figure 6: Topology Discovery Connection Phases

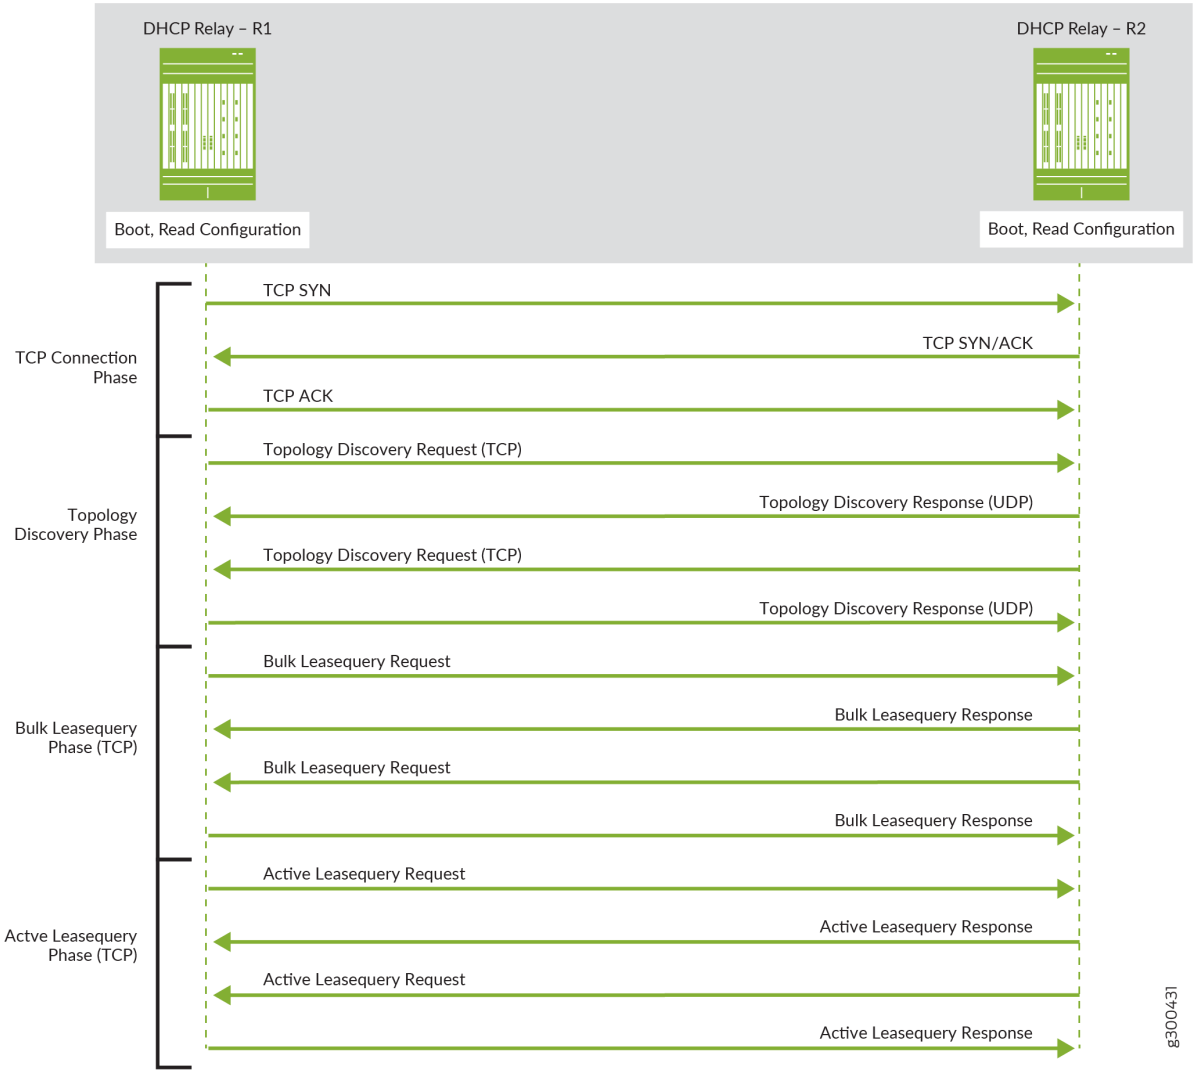

1. TCP connection phase—A TCP connection is established between the peer relay agents.
2. Topology discovery phase—The peers exchange topology discovery messages to determine the matching access interfaces for each subscriber redundancy group on the peers. The remote peer matches an interface based on 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 809 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
Vendor Specific Option (DHCPv4 option 43; DHCPv6 option 17)—This option carries the following information specific for the vendor, Juniper Networks: • Suboption 1—A string with the value topology_discover_lq. 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 giaddr 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 *logical-system-name*;*routing-instance-name*. • Suboption 6—Shared common key of the querying interface. This is an ASCII string of up to 63 characters.	Vendor Specific Option (DHCPv4 option 43; DHCPv6 option 17)—This option carries the following information: • Suboption 1—A string with the value topology_discover_lq. 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. 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. For M:N redundancy using pseudowires, matching is based on the querying interface's shared common key and transaction ID received in the request.

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 1173.

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/xe/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 811 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 751.

  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 751.

  ```
  [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 919.

  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 824.

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 829.

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 811 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 751.

  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 751.

```
[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 919.

```
[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 824.

   - 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 824.

   - 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 829.

Configuring and Using DHCP Active Leasequery

Before you begin, read "Guidelines for Configuring Support for Individual, Bulk, and Active Leasequery Operations" on page 811 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 751.

  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 751.

```
[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 919.

  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 814.

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 1173 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 824.

> **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 829.

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 825 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	–	✓	–	✓
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" on page 2628, "request dhcpv6 relay leasequery" on page 2637, "request dhcp relay bulk-leasequery" on page 2625, and "request dhcpv6 relay bulk-leasequery" on page 2634.

- 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 | 829
- Action | 829

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 830.

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 | 830
- Action | 831

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 829.

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
```

Release History Table

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 | 691
DHCPv6 Local Server | 910
DHCPv6 Relay Agent | 916

DHCP Client Authentication With An External AAA Authentication Service

IN THIS SECTION

- Specifying Authentication Support | 833
- Creating Unique Usernames for DHCP Clients | 834
- Example-Configuring DHCP with External Authentication Server | 837

Specifying Authentication Support

Include the `authentication` statement at hierarchy levels given in Table 49 on page 834. 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 (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 | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916

Receiving DHCP Options From a RADIUS Server

IN THIS SECTION

- Centrally Configure DHCP Options on a RADIUS Server | 838
- Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview | 843
- Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6 Subscribers | 846
- Monitoring DHCP Options Configured on RADIUS Servers | 849

Centrally Configure DHCP Options on a RADIUS Server

IN THIS SECTION

- RADIUS-Sourced Options | 839
- Client-Sourced Options Configuration | 840
- Data Flow for RADIUS-Sourced DHCP Options | 840
- Multiple VSA 26-55 Instances Configuration | 841

- DHCP Options That Cannot Be Centrally Configured | **842**

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 840 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

Figure 7: DHCP Options Data Flow

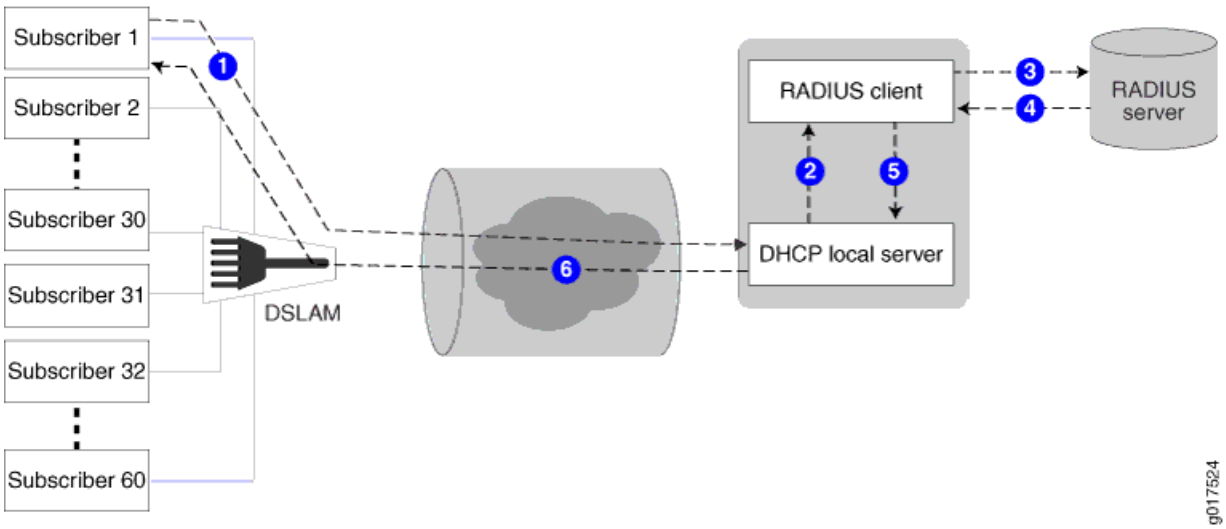

The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

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 by 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 `RO` flags in the Subscriber Management RADIUS dictionary file. The `R` value indicates a multivalued reply attribute and the `O` value indicates an ordered attribute.

DHCP Options That Cannot Be Centrally Configured

Table 50 on page 842 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.

Table 50: Unsupported Opaque DHCP Options *(Continued)*

DHCP Option	Option Name	Comments
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 | **845**
- Excluding the VSAs from RADIUS Messages | **846**

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 | **847**
- Exchange of DHCPv4 Client, DHCPv4 Server, and RADIUS-Sourced Options | **847**
- Exchange of DHCPv6 Client, DHCPv6 Server, and RADIUS-Sourced Options | **848**

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 | 849
- Action | 849
- Meaning | 850

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 is 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 851 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

Release History Table

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 | **691**
- DHCPv6 Local Server | **910**
- DHCPv6 Relay Agent | **916**
- 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 | 852
- Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces | 855
- Configuring Group-Specific DHCP Local Server Options | 856
- Configuring Group-Specific DHCP Relay Options | 857
- Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups | 858

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" on page 1944 *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 707.

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 858.

- `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 709.

- `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 727.

- 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 751.

- 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 764 .

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
```

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1
16.2R1	Starting in Junos OS Release 16.2R1

RELATED DOCUMENTATION

DHCP Overview | **691**

DHCPv6 Local Server | **910**

DHCPv6 Relay Agent | **916**

Number of DHCP Clients Per Interface

IN THIS SECTION

- Specifying the Maximum Number of DHCP Clients Per Interface | 862
- Allowing Only One DHCP Client Per Interface | 863

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 | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916
- Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706
- Conserving IP Addresses Using DHCP Auto Logout | 878

Maintaining DHCP Subscribers During Interface Delete Events

IN THIS SECTION

- Maintaining Subscribers During Interface Delete Events | 865
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | 866
- Configuring an ACX Series DHCP Local Server to Preserve Subscriber Binding Information | 867
- Verifying and Managing DHCP Subscriber Binding During Interface Delete Events | 868

Maintaining Subscribers During Interface Delete Events

IN THIS SECTION

- Benefits of Maintaining Subscriber Bindings | 866

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 | 869

- Action | 869

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 | 691

DHCPv6 Local Server | 910

DHCPv6 Relay Agent | 916

Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706

Conserving IP Addresses Using DHCP Auto Logout | 878

Dynamic Reconfiguration of Clients From a DHCP Local Server

IN THIS SECTION

- Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | 870
- Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients | 875
- Configuring Deletion of the Client When Dynamic Reconfiguration Fails | 876
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | 876
- Configuring a Token for DHCP Local Server Authentication | 877

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients

IN THIS SECTION

- Default Client/Server Interaction | 871
- Dynamic Client/Server Interaction for DHCPv4 | 871
- Dynamic Client/Server Interaction for DHCPv6 | 872
- Manually Forcing the Local Server to Initiate the Reconfiguration Process | 872
- Action Taken for Events That Occur During a Reconfiguration | 873
- Benefits of Dynamic Reconfiguration of DHCP Local Server Clients | 873

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 873 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.
The `clear dhcp server binding` command is issued.	Server deletes client.
The `request dhcp server reconfigure` (DHCPv4) or `request dhcpv6 server reconfigure` (DHCPv6) 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 renews its lease or rebind to the server.

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

The DHCP local server can initiate reconfiguration of its clients to avoid extended outages because of server configuration changes. You can enable dynamic reconfiguration for all DHCP clients or only the DHCP clients serviced by a specified group of interfaces, and you can modify the behavior accordingly.

Starting in Junos OS Release 14.1, you can modify the behavior of the process in which the DHCP local server initiates reconfiguration of its clients by including the appropriate configuration statements. You can provide the statements at the `[edit system services dhcp-local-server reconfigure]` hierarchy level for all DHCPv4 clients, and at the `[edit system services dhcp-local-server dhcpv6 reconfigure]` hierarchy level for all DHCPv6 clients. To override this global configuration for only the DHCP clients serviced by a specified group of interfaces, you can include the statements with different values at the `[edit system services dhcp-local-server group` *`group-name`* `reconfigure]` hierarchy level for DHCPv4 clients, and at the `[edit system services dhcp-local-server dhcpv6 group` *`group-name`* `reconfigure]` hierarchy level for DHCPv6 clients.

To configure dynamic reconfiguration of DHCP clients:

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 875.

4. (Optional) Configure the response to a failed reconfiguration.

   See "Configuring Deletion of the Client When Dynamic Reconfiguration Fails" on page 876.

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 876.

6. (Optional) Configure a token for rudimentary server authentication.

   See "Configuring a Token for DHCP Local Server Authentication" on page 877.

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 912.

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 forcerenew 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 forcerenew 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 forcerenew 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
```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, you can modify the behavior of the process in which the DHCP local server initiates reconfiguration of its clients by including the appropriate configuration statements.

RELATED DOCUMENTATION

- DHCP Overview | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916
- DHCP Monitoring and Management | 895

Conserving IP Addresses Using DHCP Auto Logout

IN THIS SECTION

- DHCP Auto Logout Overview | 879
- Automatically Logging Out DHCP Clients | 881
- How DHCP Relay Agent Uses Option 82 for Auto Logout | 882
- DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers | 883
- Automatically Logging Out DHCPv6 Clients | 884

DHCP Auto Logout Overview

IN THIS SECTION

- Auto Logout Overview | 879
- How DHCP Identifies and Releases Clients | 879
- Option 60 and Option 82 Requirements | 880

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 882 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
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

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”		
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 | 862
- DHCP Monitoring and Management | 895
- Overrides for Default DHCP Local Server and DHCP Relay Configuration Settings | 706
- DHCP Overview | 691
- DHCPv6 Local Server | 910
- DHCPv6 Relay Agent | 916

DHCP Short Cycle Protection

IN THIS SECTION

- DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions | 885
- Configuring DHCP Short-Cycle Protection | 889
- Verifying and Managing DHCP Short-Cycle Protection | 892

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 | 886
- How DHCP Short-Cycle Protection Works | 887
- Termination of the Lockout Condition | 888
- Benefits of Using DHCP Short Cycle Protection | 888

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:

Lockout time = (Lockout minimum time) x $[2^{(n-1)}]$

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 | 892
- Action | 893
- Meaning | 893

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
```

Release History Table

Release	Description
18.2R1	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.

RELATED DOCUMENTATION

DHCP Overview | 691

DHCPv6 Local Server | 910

DHCPv6 Relay Agent | 916

DHCP Monitoring and Management

IN THIS SECTION

- Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings | 895
- Viewing and Clearing DHCP Bindings | 896
- Monitoring DHCP Relay Server Responsiveness | 898
- Verifying DHCP Server Binding and Server Statistics | 899
- Verifying and Managing DHCP Relay Configuration | 901
- Tracing Extended DHCP Operations | 902

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 | 899
- Action | 900

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 | 901
- Action | 901

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 | **904**
- Configuring the Number and Size of Extended DHCP Log Files | **904**
- Configuring Access to the Extended DHCP Log File | **905**
- Configuring a Regular Expression for Extended DHCP Messages to Be Logged | **906**
- Configuring the Extended DHCP Tracing Flags | **906**
- Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged | **907**
- Tracing Extended DHCP Operations for Specific Interfaces | **908**

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 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 902.

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 902.

   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 902.

      - Configure the number and size of the log files.

        See "Configuring the Number and Size of Extended DHCP Log Files" on page 902.

      - Configure access to the log file.

        See "Configuring Access to the Extended DHCP Log File" on page 902.

      - Configure a regular expression to filter logging events.

See "Configuring a Regular Expression for Extended DHCP Messages to Be Logged" on page 902.

c. (Optional) Specify tracing flag options.

See "Configuring the Extended DHCP Tracing Flags" on page 902.

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 902.

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 | 870

DHCPv6 Monitoring and Management | 925

Conserving IP Addresses Using DHCP Auto Logout | 878

DHCP Overview | 691

DHCPv6 Local Server | 910

DHCPv6 Relay Agent | 916

CHAPTER 6

DHCPv6 for Subscriber Management

IN THIS CHAPTER

- DHCPv6 Local Server | **910**
- DHCPv6 Relay Agent | **916**
- DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | **923**
- DHCPv6 Monitoring and Management | **925**

DHCPv6 Local Server

IN THIS SECTION

- DHCPv6 Local Server Overview | **910**
- Configuring DHCPv6 Rapid Commit (MX Series, EX Series) | **912**
- Preventing Binding of Clients That Do Not Support Reconfigure Messages | **912**
- Configuring the DUID Type Supported by DHCPv6 Servers | **913**
- Example: Extended DHCPv6 Local Server Configuration | **914**

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 54 on page 911 to configure the client:

Table 54: 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 | 707

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 | **874**
>
> Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients | **870**

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 | **691**

DHCPv6 Relay Agent | **916**

DHCPv6 Relay Agent

IN THIS SECTION

- DHCPv6 Relay Agent Overview | **916**
- DHCPv6 Relay Agent Options | **917**
- Configuring DHCPv6 Relay Agent Options | **917**
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | **919**
- Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets | **921**

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 925 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" on page 2278 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 | 751

Using DHCP Relay Agent Option 82 Information | 751

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 | 751

Using DHCP Relay Agent Option 82 Information | 751

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 | 762

Release History Table

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 | 910

DHCP Overview | 691

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking

IN THIS SECTION

- Benefits of Client MAC address Validation | **924**

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.

Release History Table

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 | 910

DHCPv6 Relay Agent Overview | 916

DHCPv6 Monitoring and Management

IN THIS SECTION

- Requesting DHCPv6 Local Server to Initiate Reconfiguration of Client Bindings | 925
- Verifying and Managing DHCPv6 Local Server Configuration | 927
- Verifying and Managing DHCPv6 Relay Configuration | 928

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 | 870

Verifying and Managing DHCPv6 Local Server Configuration

IN THIS SECTION

- Purpose | 927
- Action | 927

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 | 928
- Action | 928

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 | 910

DHCPv6 Relay Agent | 916

3

PART

IPv6 for Subscriber Management

IPv6 for Subscriber Management | 930

CHAPTER 7

IPv6 for Subscriber Management

IN THIS CHAPTER

- Introduction to IPv6 Addresses | **930**
- Migration to IPv6 Using IPv4 and IPv6 Dual Stack | **934**
- IPv6 WAN Link Addressing with NDRA | **939**
- IPv6 WAN Link Addressing with DHCPv6 IA_NA | **946**
- Subscriber LAN Addressing with DHCPv6 Prefix Delegation | **948**
- WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **956**
- Designs for IPv6 Addressing in a Subscriber Access Network | **993**
- Dual-Stack Access Models in a DHCP Network | **1000**
- Dual-Stack Access Models in a PPPoE Network | **1012**
- Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | **1039**
- Dual Stack for PPPoE Access Networks Using DHCP | **1044**
- Dual Stack for PPPoE Access Networks Using NDRA | **1048**
- Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1114**
- Dual Stack Subscribers Monitoring and Management | **1125**

Introduction to IPv6 Addresses

IN THIS SECTION

- IPv6 Notation | **931**
- IPv6 Prefixes | **931**
- IPv6 Address Types | **932**

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:fe0.

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 | 932
- Multicast Addresses | 933
- Anycast Addresses | 933

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 | **937**

Migration to IPv6 Using IPv4 and IPv6 Dual Stack

IN THIS SECTION

- Basic Architecture of a Subscriber Access Dual-Stack Network | 934
- Terms Used in IPv6 Subscriber Management Documentation | 935
- IPv6 Addressing Requirements for a Subscriber Access Network | 937

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 935 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

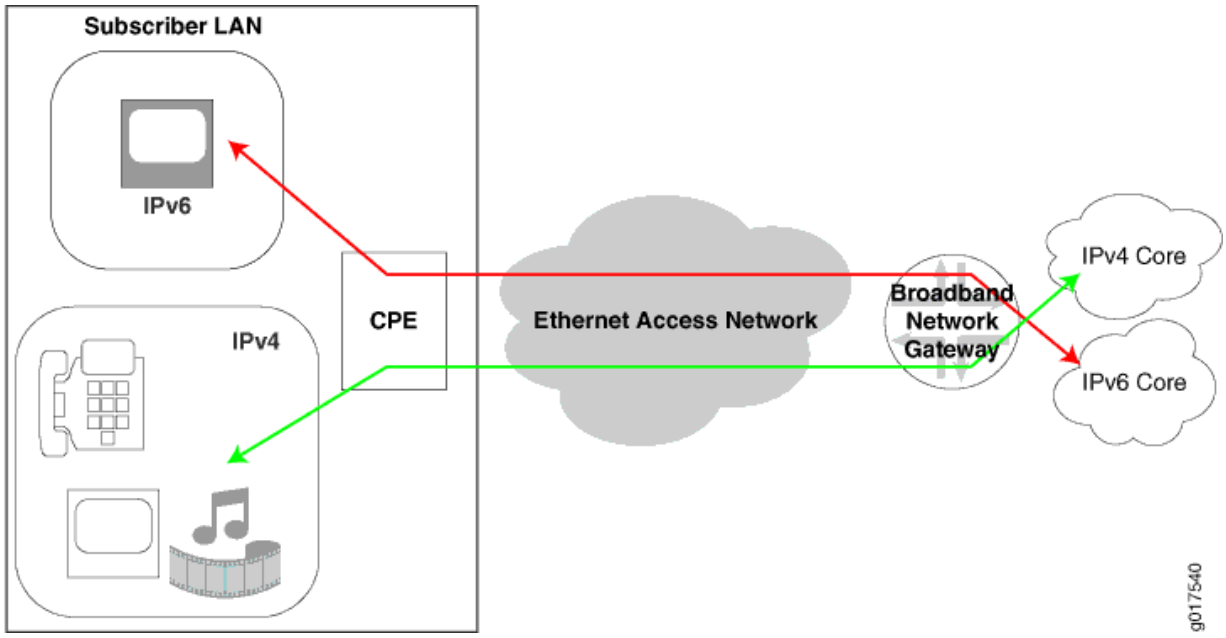

Terms Used in IPv6 Subscriber Management Documentation

Table 55 on page 935 defines terms used in the IPv6 subscriber management documentation.

Table 55: 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 55: IPv6 Subscriber Management Terms *(Continued)*

Term	Definition
DHCPv6 IA	Identity association. A collection of addresses assigned to a client. Each IA contains one type of address. For example, IA_NA carries assigned addresses that are nontemporary addresses; IA_PD carries a prefix.
DHCPv6 IA_PD	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. Referred to as DHCPv6 prefix delegation.
DHCPv6 IA_NA	IA for nontemporary addresses. An IA that carries assigned addresses that are not temporary addresses. DHCPv6 IA_NA is used to assign global IPv6 addresses.
Global IPv6 address	Unique IPv6 address that identifies a single interface and allows the interface to access the IPv6 internet.
IPv6 address prefix/ prefix length	Combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form *ipv6-prefix/prefix-length* and represents a block of address space (or a network). The */prefix-length* indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 2001:DB8::/32 is an IPv6 prefix.
IPCP	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.
IPv6CP	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.

Table 55: IPv6 Subscriber Management Terms *(Continued)*

Term	Definition
Link-local address	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. In your dual-stack network, you can use a link-local address on the interface that connects the CPE and the BNG.
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)	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. In a dual-stack network, the router sends RAs to CPE devices on its access network.
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 | **938**

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 938 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

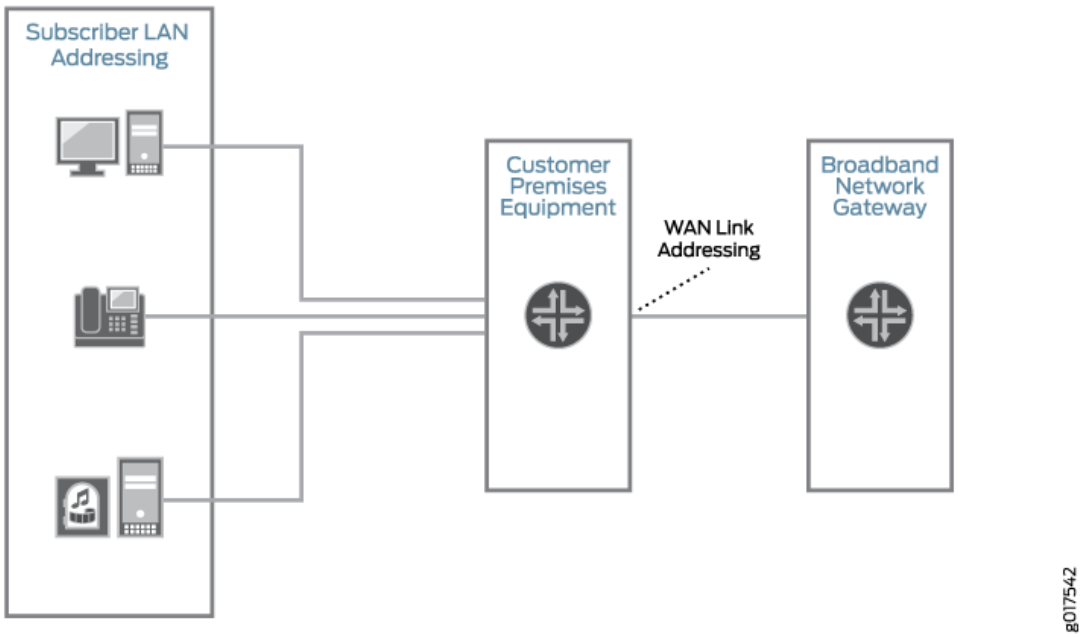

You can use the following methods for assigning IPv6 addresses:

- 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 | 939
IPv6 WAN Link Addressing with DHCPv6 IA_NA | 946
Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 948
WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | 956

IPv6 WAN Link Addressing with NDRA

IN THIS SECTION

Using NDRA to Provide IPv6 WAN Link Addressing Overview | 939
IPv6 Neighbor Discovery Protocol Overview | 941
Dynamic Router Advertisement Configuration Overview | 942
Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | 942
Methods for Obtaining IPv6 Prefixes for NDRA | 944
Duplicate Prefix Protection for NDRA | 945

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 | **997**

Design 3: IPv6 Addressing with NDRA | **998**

IPv6 Neighbor Discovery Protocol Overview

IN THIS SECTION

- Neighbor Discovery Messages | **941**

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 | **1050**

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 | **945**

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 | **1050**

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 | **1051**

Release History Table

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 | **934**
- Designs for IPv6 Addressing in a Subscriber Access Network | **993**
- Dual-Stack Access Models in a PPPoE Network | **1012**
- Dual Stack for PPPoE Access Networks Using NDRA | **1048**

IPv6 WAN Link Addressing with DHCPv6 IA_NA

IN THIS SECTION

- Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA | **947**
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | **947**

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 | **947**

You can set up the BNG to select global IPv6 addresses 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 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 | **934**
- Designs for IPv6 Addressing in a Subscriber Access Network | **993**
- Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | **1039**
- Dual-Stack Access Models in a PPPoE Network | **1012**
- Dual Stack for PPPoE Access Networks Using DHCP | **1044**

Subscriber LAN Addressing with DHCPv6 Prefix Delegation

IN THIS SECTION

- Using DHCPv6 Prefix Delegation Overview | **949**
- Using a Delegated Prefix on the CPE Loopback Interface | **950**
- DHCPv6 Prefix Delegation over PPPoE | **950**
- Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation | **951**
- DHCPv6 Prefix Exclusion | **952**
- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation | **954**
- Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | **955**

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 949 shows how DHCPv6 prefix delegation is used in a dual-stack network.

Figure 10: Delegated Addressing in a Dual-Stack Network Using DHCPv6

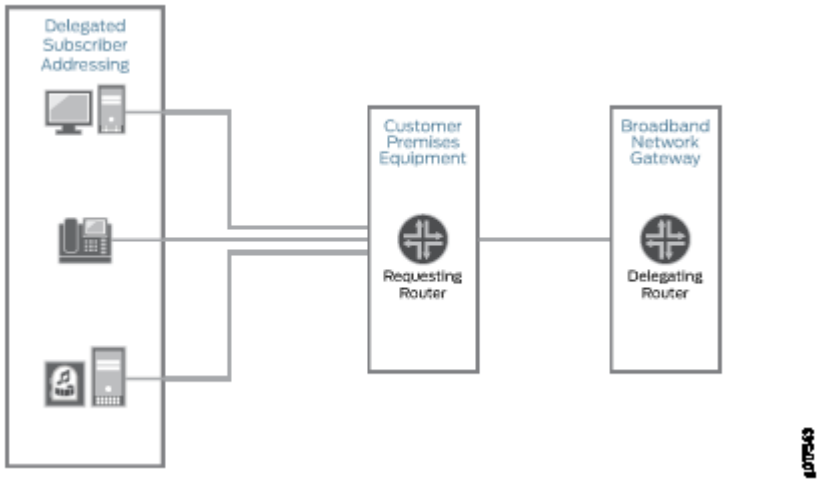

DHCPv6 prefix delegation operates as follows:

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 | **996**
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | **997**
- Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | **999**
- Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | **994**

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 | **993**

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 | **1078**

Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation

IN THIS SECTION

- Using a AAA RADIUS Server to Obtain IPv6 Prefixes for Prefix Delegation | **951**

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 | 995

DHCPv6 Prefix Exclusion

IN THIS SECTION

- Configuring DHCPv6 Prefix Exclude Option | 953

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 952 shows how DHCPv6 prefix delegation is used in a dual-stack network.

Figure 11: Delegated Addressing in a Dual-Stack Network Using DHCPv6

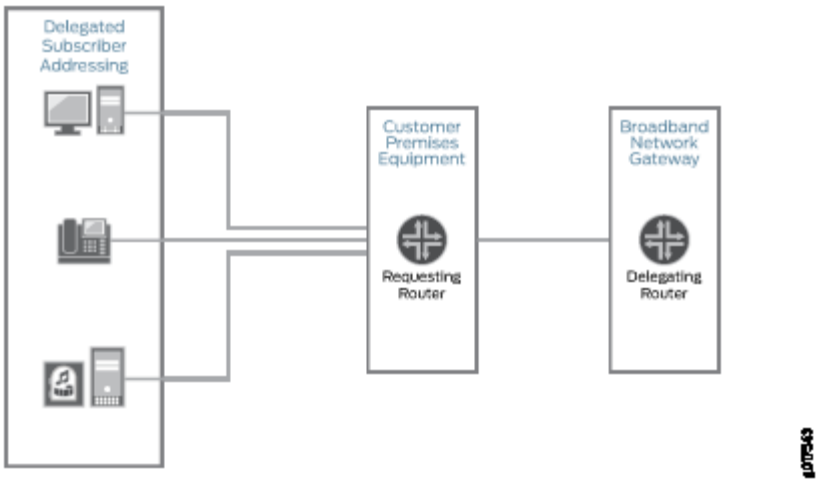

DHCPv6 prefix delegation operates as follows:

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-assignement 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 | 934

Designs for IPv6 Addressing in a Subscriber Access Network | 993

WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation

IN THIS SECTION

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 957
- DHCPv6 Options in a DHCPv6 Multiple Address Environment | 958
- Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA | 959
- Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface | 961
- Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | 961

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 | **957**
- Behavior When CPE Sends Separate Renew Requests for IA_NA and IA_PD Address Types | **957**

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 | **946**

Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **996**

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 | **957**

Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | **1169**

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 | **960**
- Using a AAA RADIUS Server to Obtain IPv6 Addresses and Prefixes | **960**
- Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment | **960**

You can set up the BNG to select global IPv6 addresses 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 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 | 954

Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA | 947

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 | 961
- Overview | 962
- Configuration | 964
- Verification | 986

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 | 962

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

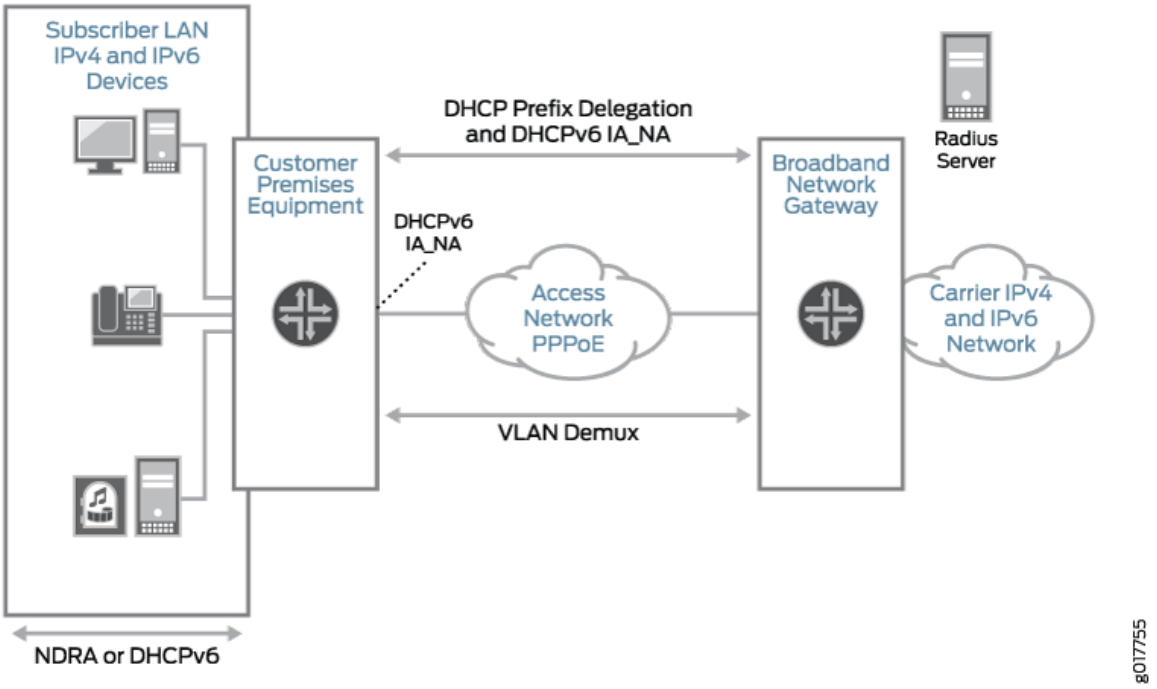

Table 56 on page 963 describes the configuration components used in this example.

Table 56: 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.
	demux0	VLAN demux interface that runs over the underlying Ethernet interface.

Table 56: Configuration Components Used in Dual Stack with DHCPv6 IA_NA and DHCPv6 Prefix Delegation *(Continued)*

Configuration Component	Component Name	Purpose
	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 | **965**
- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | **968**
- Configuring a Dynamic Profile for the PPPoE Logical Interface | **970**
- Configuring a Loopback Interface | **973**
- Configuring a VLAN Demux Interface over an Ethernet Underlying Interface | **975**
- Configuring an Interface for Communication with RADIUS Server | **977**
- Specifying the BNG IP Address | **978**
- Configuring RADIUS Server Access | **980**
- Configuring RADIUS Server Access Profile | **982**
- Configuring Local Address-Assignment Pools | **983**

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;
                        preferred;
                    }
                }
                family inet6 {
                    address 2001:db8:0::1/128 {
```

```
                    primary;
                    preferred;
                }
            }
        }
    }
}
routing-options {
    router-id 203.0.113.0;
}
access {
    radius-server {
        203.0.113.99 {
            secret "$ABC123$ABC123$ABC123"; ## 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 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 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 | **986**
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | **987**
- Verifying Dynamic Subscriber Sessions | **988**
- Verifying DHCPv6 Address Pools Used for DHCPv6 Prefix Delegation | **989**
- Verifying DHCPv6 Address Bindings | **990**
- Verifying PPP Options Negotiated with the Remote Peer | **991**

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 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 | **996**

Release History Table

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 | **946**

Subscriber LAN Addressing with DHCPv6 Prefix Delegation | **948**

Designs for IPv6 Addressing in a Subscriber Access Network | **993**

Designs for IPv6 Addressing in a Subscriber Access Network

IN THIS SECTION

- Selecting the Type of Addressing Used on the CPE | **993**
- Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link | **993**
- Selecting the Method of Assigning Global IPv6 Addresses to Subscribers | **994**
- Selecting the Method of Obtaining IPv6 Prefixes | **995**
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **996**
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation | **997**
- Design 3: IPv6 Addressing with NDRA | **998**
- Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix | **999**

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 | **937**

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 57 on page 994.

Table 57: 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 | 937

Using NDRA to Provide IPv6 WAN Link Addressing Overview | 939

IPv6 WAN Link Addressing with DHCPv6 IA_NA | 946

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 | 949

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | 957

Selecting the Method of Obtaining IPv6 Prefixes

IN THIS SECTION

Using a AAA RADIUS Server to Obtain Global IPv6 Addresses and IPv6 Prefixes | 995

Using a Local Pool to Assign IPv6 Addresses or Prefixes | 996

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 58 on page 995 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 58: 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 58: 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 | **937**

Using DHCPv6 Prefix Delegation Overview | **949**

Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

This design (Figure 13 on page 997) 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

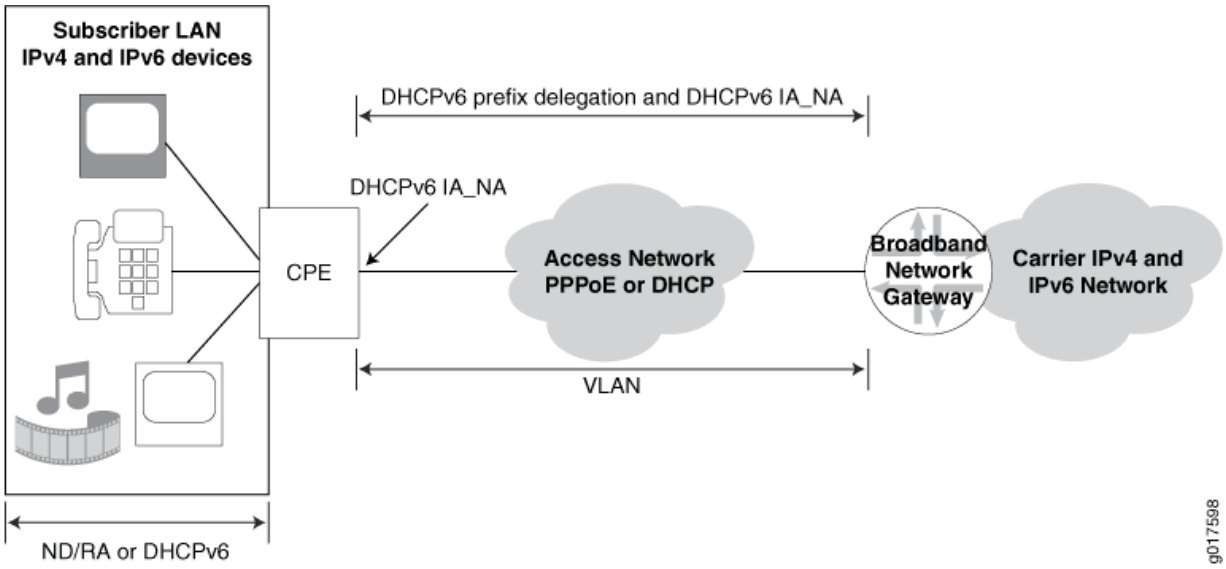

SEE ALSO

Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE | **961**

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview | **957**

Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation

This design (Figure 14 on page 998) 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

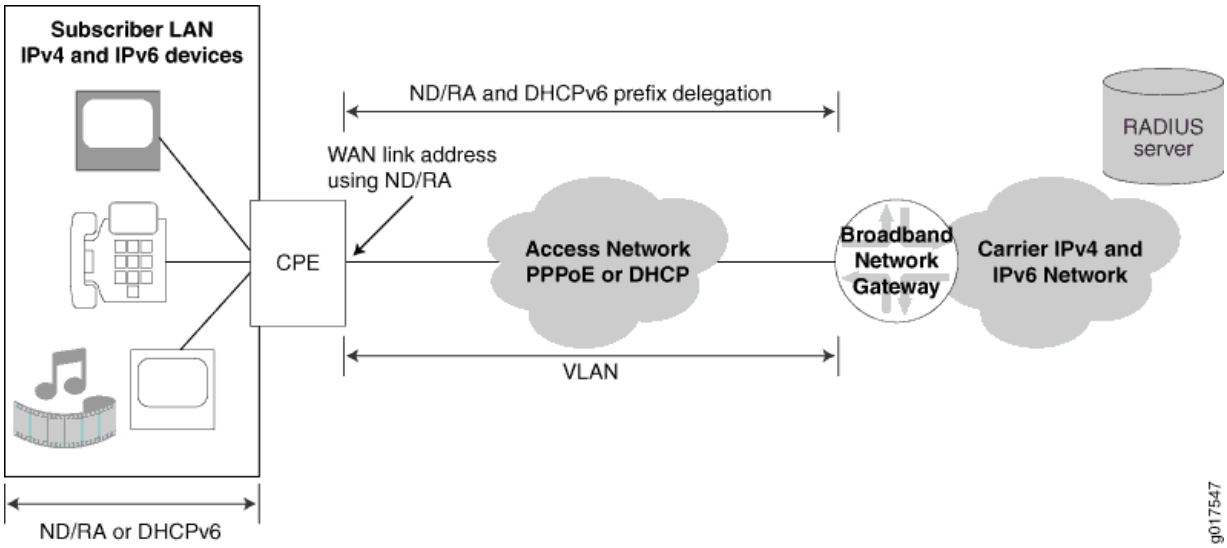

If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of design 2 and design 3.

SEE ALSO

Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | **1078**

Using DHCPv6 Prefix Delegation Overview | **949**

Design 3: IPv6 Addressing with NDRA

In this design (Figure 15 on page 999), 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

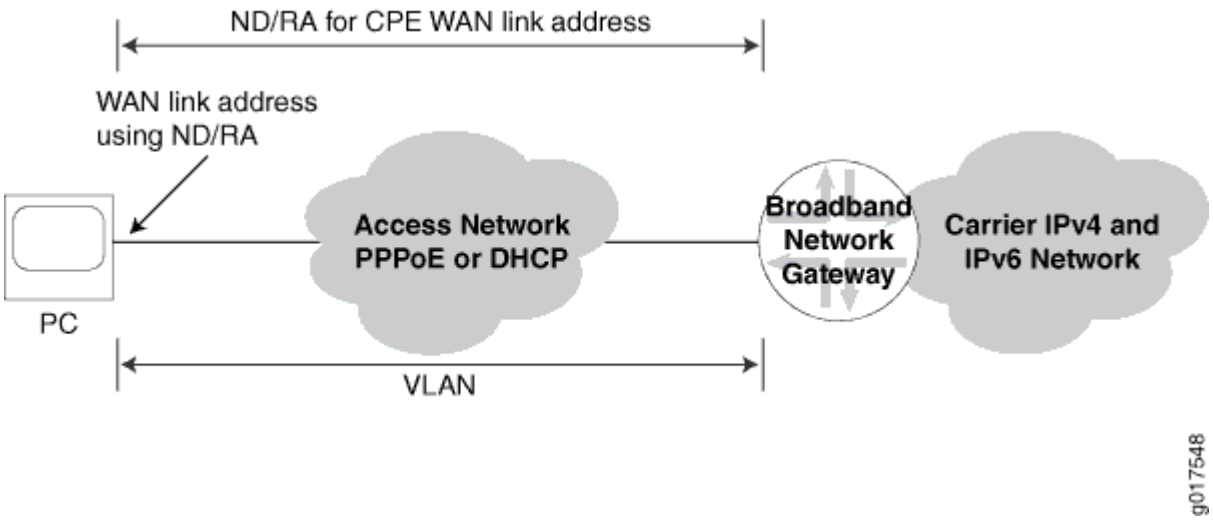

If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of Design 2 and Design 3.

SEE ALSO

Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | **1052**

Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix

In this design (Figure 16 on page 1000), 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

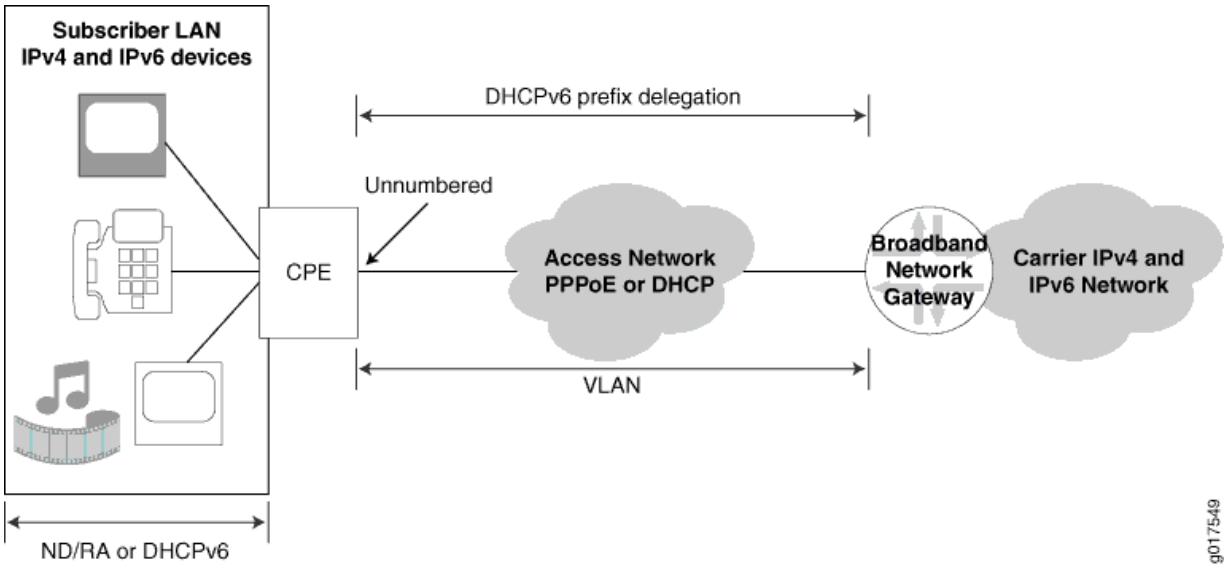

SEE ALSO

- Using DHCPv6 Prefix Delegation Overview | **949**

RELATED DOCUMENTATION

- Migration to IPv6 Using IPv4 and IPv6 Dual Stack | **934**
- Subscriber LAN Addressing with DHCPv6 Prefix Delegation | **948**
- IPv6 WAN Link Addressing with NDRA | **939**
- Dual Stack for PPPoE Access Networks Using NDRA | **1048**
- WAN and LAN Addressing Using DHCPv6 IA_NA and DHCPv6 Prefix Delegation | **956**

Dual-Stack Access Models in a DHCP Network

IN THIS SECTION

- IPv4 and IPv6 Dual Stack in a DHCP Access Network | **1001**

- AAA Service Framework in a Dual Stack over a DHCP Access Network | **1002**
- Dual-Stack Interface Stack in a DHCP Wholesale Network | **1003**
- Single-Session DHCP Dual-Stack Overview | **1004**
- Configuring Single-Session DHCP Dual-Stack Support | **1008**
- Verifying and Managing DHCP Dual-Stack Configuration | **1011**

IPv4 and IPv6 Dual Stack in a DHCP Access Network

IN THIS SECTION

- Support for Demultiplexing Interfaces | **1002**

Figure 17 on page 1001 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

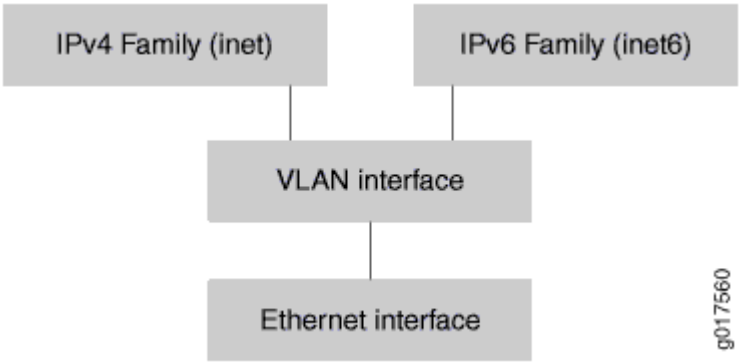

NOTE: When you are using IPv4 and IPv6 dual stack on the same DHCP interface, you must configure one dynamic profile for both the IPv4 and IPv6 subscribers. You cannot run IPv4 and IPv6 subscriber sessions over the same interface if you configure separate dynamic profiles for IPv4 and IPv6.

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 | 934

AAA Service Framework in a Dual Stack over a DHCP Access Network

IN THIS SECTION

- Collection of Accounting Statistics in a DHCP Access Network | 1003
- Change of Authorization (CoA) | 1003

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 1003, 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

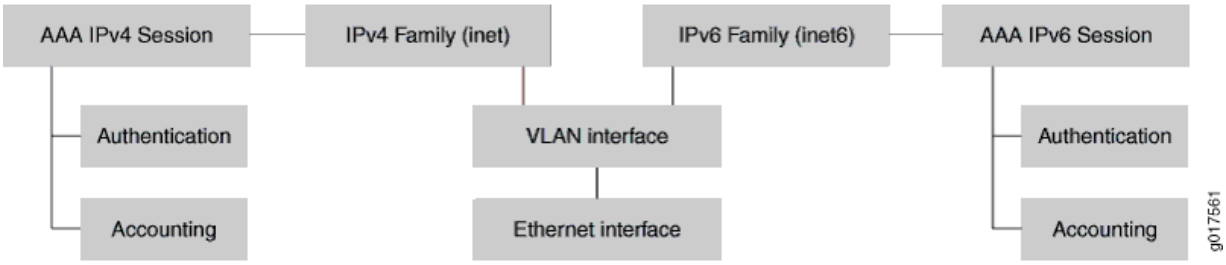

Collection of Accounting Statistics in a DHCP Access Network

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 1004 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

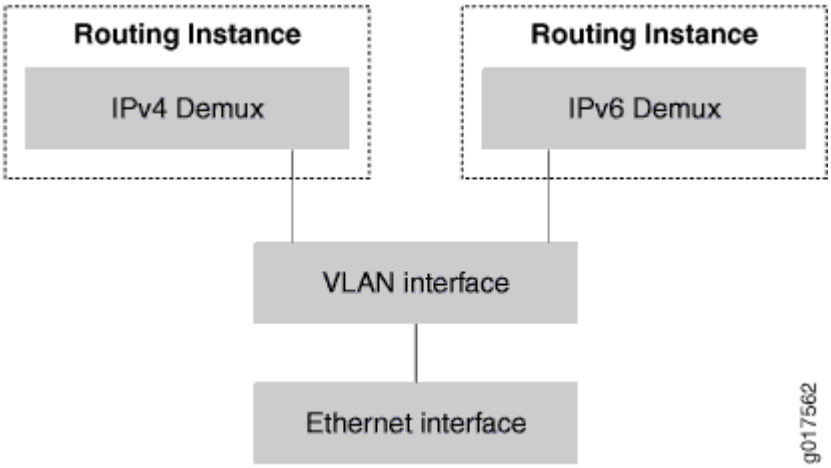

Single-Session DHCP Dual-Stack Overview

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 legs 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 1005, single subscriber session is established for dual-stack user.

Figure 20: DHCP Dual Stack Single-Session Subscriber Deployment Model

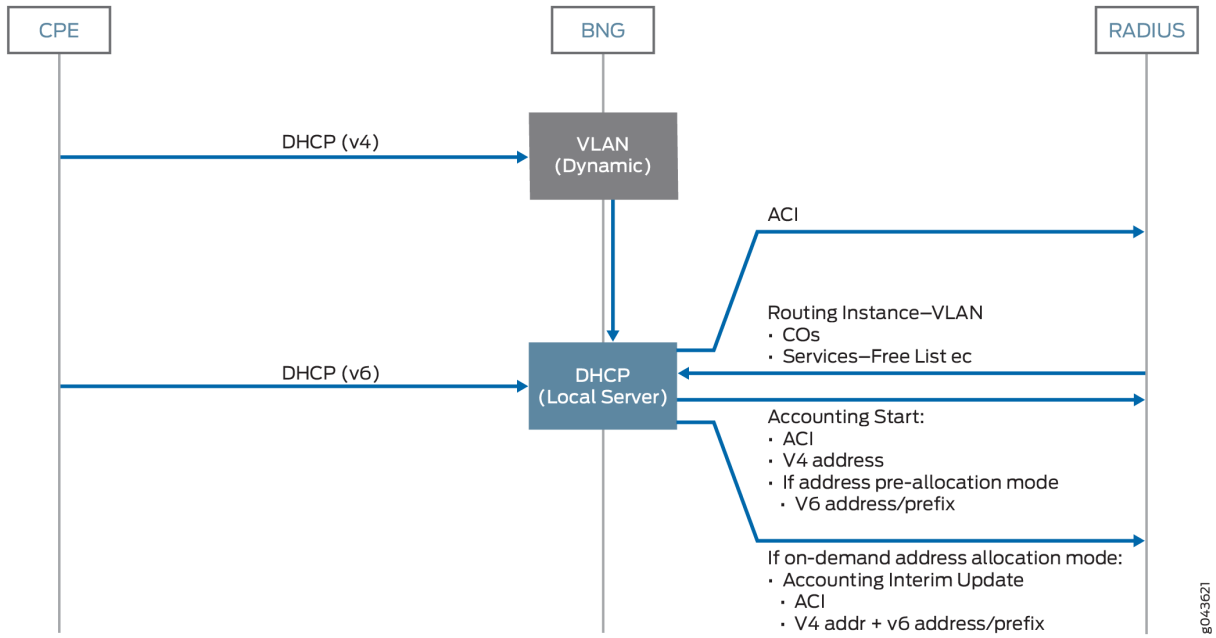

You can configure single-session dual-stack subscriber settings for DHCP relay agent and DHCP local server. You use the `dual-stack-group` statement to create a named group that specifies the values for dual stack subscribers. Then, you use the `dual-stack` statement to specify the name of the dual stack 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 59 on page 1007 describes the behavior.

Table 59: 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 | 556

Configuring RADIUS Reauthentication for DHCP Subscribers | 567

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 703.

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 833.

- 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 834.

- 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 837.

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 919.

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 751.

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 921.

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 751.

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 709.

   - 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 709.

10. (Optional) Verify your dual-stack group configuration for DHCPv4 and DHCPv6.

    See "Verifying and Managing DHCP Dual-Stack Configuration" on page 1011.

```
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 | 1011
- Action | 1011

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.

Release History Table

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 | 934

Dual-Stack Access Models in a DHCP Network | 1000

Dual-Stack Access Models in a PPPoE Network

IN THIS SECTION

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1013
- Shared IPv4 and IPv6 Service Sessions on PPP Access Networks | 1016
- AAA Service Framework in a Dual Stack over a PPPoE Access Network | 1016
- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers | 1019
- Accounting Messages for PPPoE Using NDRA Prefixes | 1020
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes | 1027
- Suppressing Accounting Information That Comes from AAA | 1037
- Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address | 1038

IPv4 and IPv6 Dual Stack in a PPPoE Access Network

IN THIS SECTION

- Support for Demultiplexing Interfaces | **1015**
- Determining the Status of CPE in a PPPoE Access Network | **1015**
- IPV6 Address Provisioning in the PPPoE Access Network | **1015**
- Authentication in a PPPoE Access Network | **1015**
- Negotiation of Network Control Protocols When Authorized Addresses Are Unavailable | **1015**

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 1014 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

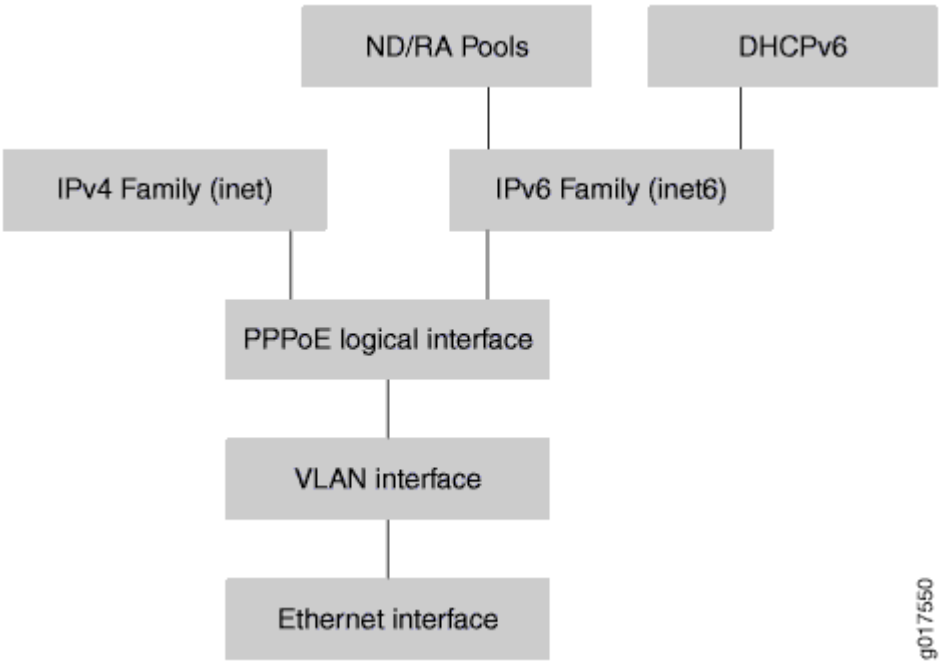

Figure 22 on page 1014 shows a dual-stack interface stack over aggregated Ethernet in a PPPoE access network.

Figure 22: Dual-Stack Aggregated Ethernet Stack over a PPPoE Access Network

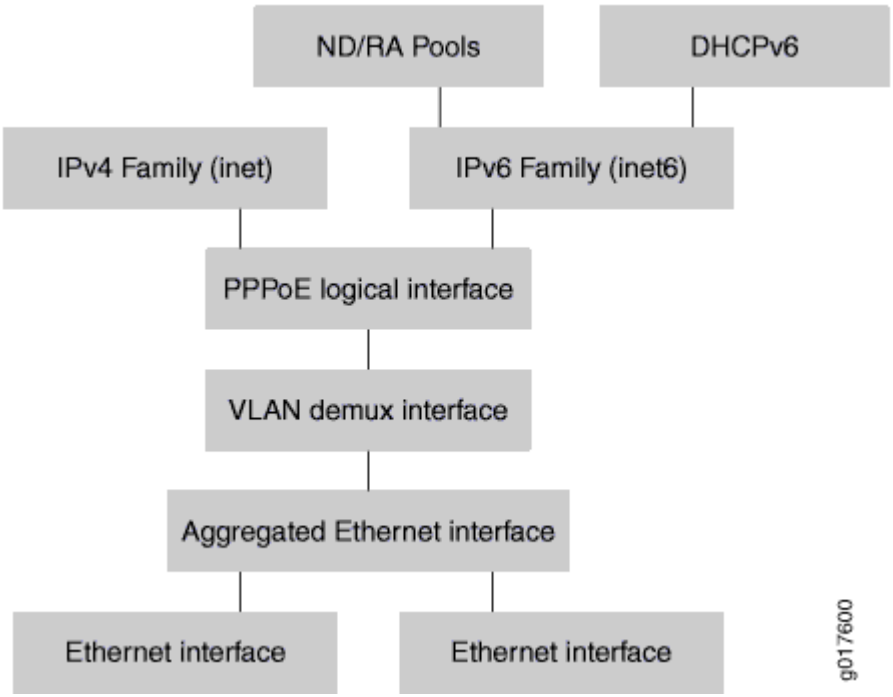

Support for Demultiplexing Interfaces

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 | **1016**
- Deactivating Shared IPv4 and IPv6 Service Sessions | **1016**

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 | **1018**

Change of Authorization (CoA) | 1019

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 1018, 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

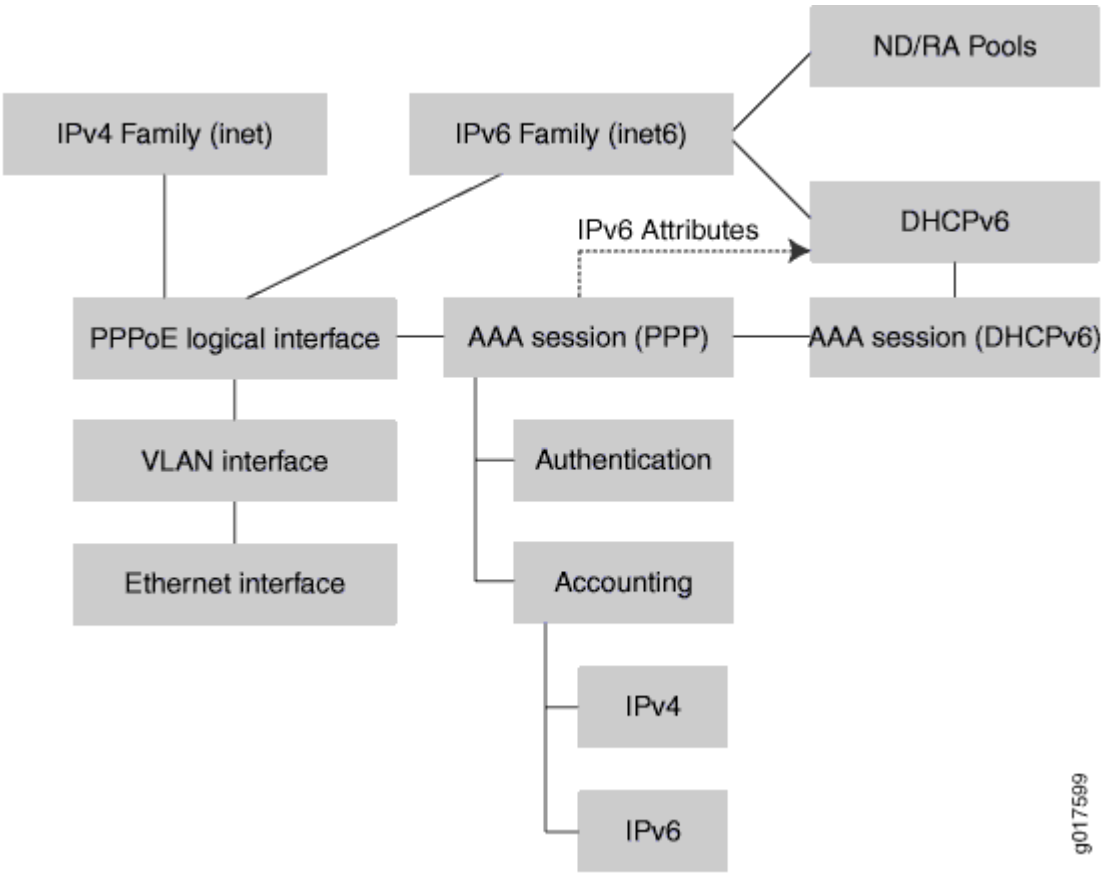

Collection of Accounting Statistics in a PPPoE Access Network

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 Solict, 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	Framed-IPv6-Prefix (used for NDRA Prefix) Framed-Interface-Id not sent Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix and Delegated-IPv6-Prefix. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6. Next periodic Acct-Interim-Update message (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.

The following table lists prefixes from RADIUS selected pools:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id Jnpr-Delegated-IPv6-pool (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id. Framed-IPv6-Prefix is based on the configuration present in the dynamic profile IPv6 prefix that was allocated and sent in Acct-Start message. Framed-IPv6-Pool is learned from RADIUS. Delegated-IPv6-Prefix is pre-allocated. Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS. No Acct-Interim-Update message is sent.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id. Framed-IPv6-Prefix and Framed-IPv6-Prefix are based on the configuration present in dynamic profile IPv6 prefix and is allocated prior and sent in Acct-Start message. Delegated-IPv6-Prefix is pre-allocated. Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. Delegated-IPv6-Prefix is pre-allocated. No immediate Acct-Interim-Update message is sent. 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.)
4	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. Delegated-IPv6-Prefix is pre-allocated. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6. 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.)

The following table lists prefixes from a local pool or an external server:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool. No immediate Acct-Interim-Update message is sent. Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Framed-IPv6-Pool.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	Framed-IPv6-Pool (used for NDRA Prefix) 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. Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent [IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool. No immediate Acct-Interim-Update message is sent upon DHCPv6. Upon DHCPv6, immediate Acct-Interim-Update is sent which contains Framed-IPv6-Pool, Framed-IPv6-Prefix, Framed-Interface-Id, and DHCPv6 IA_PD Prefix. Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer. DHCPv6 IA_PD is learned by DHCPv6 either by DHCPv6 external server or reservation from a local pool during DHCPv6 SARR phase. 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.)

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	Framed-IPv6-Prefix (used for IA_NA prefix) Framed-Interface-Id Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP	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.
2	Framed-IPv6-Prefix (used for IA_NA Prefix) Framed-Interface-Id Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	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. There is no immediate Acct-Interim-Update message sent after DHCPv6.
3	Framed-IPv6-Prefix (used for IA_NA Prefix) Framed-Interface-Id not sent Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message message contains Framed-IPv6-Prefix and Delegated-IPv6-Prefix. No immediate Acct-Interim-Update message is sent. Next periodic Acct-Interim-Update (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	Framed-IPv6-Prefix (used for IA_NA Prefix) Framed-Interface-Id not sent Delegated-IPv6-Prefix (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message message contains iFramed-IPv6-Prefix and Delegated-IPv6-Prefix. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6. Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id in addition to Framed-IPv6-Prefix and Delegated-IPv6-Prefix.

The following table lists prefixes from RADIUS selected pools:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	Framed-IPv6-Pool (used for IA_NA Prefix) Framed-Interface-Id Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id . Framed-IPv6 Prefix is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. Delegated-IPv6-Prefix is pre-allocated . Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS. No Acct-Interim-Update message is sent.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	Framed-IPv6-Pool (used for IA_NAPrefix) Framed-Interface-Id Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Delegated-IPv6-Prefix, and Framed-Interface-Id. Framed-IPv6 is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. Delegated-IPv6-Prefix is pre-allocated. Framed-Interface-Id is sent in Acct-Start message because it is learned from RADIUS. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6.

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	Framed-IPv6-Pool (used for IA_NA Prefix) Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. Framed-IPv6-Prefix is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. Delegated-IPv6-Prefix is pre-allocated. No immediate Acct-Interim-Update message is sent. 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.)

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	Framed-IPv6-Pool (used for IA_NA Prefix) Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool (used for DHCPv6 IA_PD)	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix, Framed-IPv6-Pool, and Delegated-IPv6-Prefix. Framed-IPv6-Prefix is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. Delegated-IPv6-Prefix is pre-allocated. No immediate Acct-Interim-Update message is sent upon IPv6NCP. No immediate Acct-Interim-Update message is sent upon DHCPv6. 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.)

The following table lists prefixes from a local pool or an external server:

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
1	Framed-IPv6-Pool (used for IA_NA Prefix) Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent	IPv6NCP	Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool. Framed-IPv6-Prefix is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. No immediate Acct-Interim-Update message is sent. 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.)

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
2	Framed-IPv6-Pool (used for IA_NA Prefix) Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent [IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]	IPv6NCP + DHCPv6	Acct-Start message contains Framed-IPv6-Prefix and Framed-IPv6-Pool. Framed-IPv6-Prefix is pre-allocated. Framed-IPv6-Pool is learned from RADIUS. No immediate Acct-Interim-Update message is sent upon DHCPv6. 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. Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer. DHCPv6 IA_PD is learned by DHCPv6 either by DHCPv6 External Server or reservation from a local pool during DHCPv6 SARR phase. 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.)

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
3	Framed-IPv6-Pool not sent [IA_NA Prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6] Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent [IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]	IPv6NCP	Acct-Start message does not contain any of the IPv6-related prefixes and addresses or pool names. No immediate Acct-Interim-Update message is sent upon IPv6NCP. Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id. (This value is learned during IPv6NCP negotiation with the peer.)

(Continued)

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
4	Framed-IPv6-Pool not sent [IA_NA Prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6] Framed-Interface-Id not sent Jnpr-Delegated-IPv6-Pool not sent [IA_PD prefix is learned from DHCPv6 External Server (DHCPv6 Relay / Relay Proxy model) or reservation from a local pool by DHCPv6]	IPv6NCP+DHCPv6	Acct-Start message does not contain any of the IPv6-related prefixes and addresses or pool names. No immediate Acct-Interim-Update message is sent upon IPv6NCP. Next periodic Acct-Interim-Update (based on interval) contains Framed-Interface-Id. (This value is learned during IPv6NCP negotiation with the peer.) Upon DHCPv6, immediate Acct-Interim-Update is sent which contains Framed-IPv6-Prefix, Framed-IPv6-Pool, Framed-Interface-Id, and DHCPv6 IA_PD. Framed-IPv6-Prefix is the IA_NA prefix learned by DHCPv6 (either by external server or reservation from a local pool). Framed-IPv6-Pool is sent only if there is a reservation of an IA_NA prefix from local pool by DHCPv6. Framed-Interface-Id value is learned during IPv6NCP negotiation with the peer. DHCPv6 IA_PD prefix is learned by DHCPv6 either by DHCPv6 external server or reservation from a local pool during DHCPv6 SARR phase. Any periodic Acct-Interim-Update before DHCPv6 completion contains Framed-Interface-Id only.

(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 | **934**
- Dual Stack for PPPoE Access Networks Using DHCP | **1044**
- Dual Stack for PPPoE Access Networks Using NDRA | **1048**
- Conservation of IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1114**

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

IN THIS SECTION

- Best Practice: Static PPPoE Interfaces with NDRA | **1039**
- Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network | **1040**
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | **1041**
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 | **1041**
- Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles | **1042**
- Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | **1044**

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 lo0.0;
                }
            }
        }
    }
}
```

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 lo0.0;
                }
            }
        }
    }
}
```

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 | **1012**
- Dual Stack for PPPoE Access Networks Using DHCP | **1044**
- Dual Stack for PPPoE Access Networks Using NDRA | **1048**

Dual Stack for PPPoE Access Networks Using DHCP

IN THIS SECTION

- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | **1045**
- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network | **1046**

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 | 1013

Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network | 1044

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 | 1041

RELATED DOCUMENTATION

Dual-Stack Access Models in a PPPoE Network | 1012

Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1039

Dual Stack for PPPoE Access Networks Using NDRA

IN THIS SECTION

- Configuring a Static PPPoE Logical Interface for NDRA | 1049
- Configuring an Address-Assignment Pool Used for Router Advertisements | 1050
- Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | 1051
- Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | 1052
- Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE | 1052

- Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE | **1078**
- IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | **1110**
- Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | **1111**

Configuring a Static PPPoE Logical Interface for NDRA

To configure a static PPPoE logical interface for static 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 | **1039**

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 | 939

Methods for Obtaining IPv6 Prefixes for NDRA | 944

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 | 939

Duplicate Prefix Protection for NDRA | 945

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 | 1053

- Overview | **1053**
- Configuration | **1055**
- Verification | **1073**

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 | **1054**

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

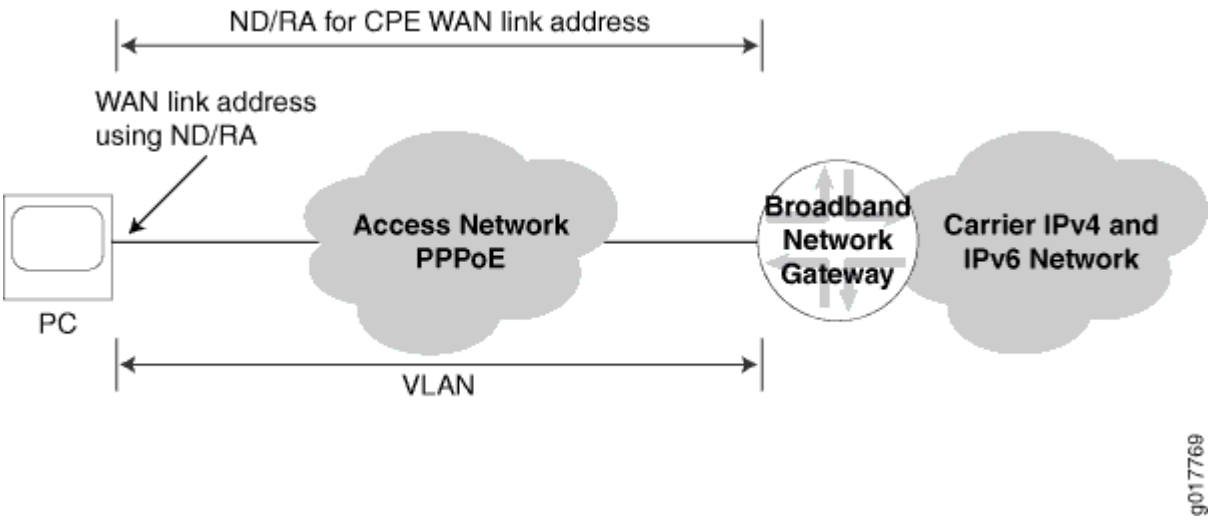

Table 60 on page 1054 describes the configuration components used in this example.

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.
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 | **1055**
- Configuring a Dynamic Profile for the PPPoE Logical Interface | **1058**
- Configuring a Loopback Interface | **1062**
- Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces | **1063**
- Specifying the BNG IP Address | **1065**
- Configuring RADIUS Server Access | **1066**
- Configuring RADIUS Server Access Profile | **1068**
- Specifying the RADIUS Server Access Profile to Use | **1070**
- Configuring Local Address-Assignment Pools | **1071**

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 | 1073
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | 1074
- Verifying Dynamic Subscriber Sessions | 1074
- Verifying the ND/RA Prefix Pool and Prefix Length | 1075
- Verifying the Status of the PPPoE Logical Interface | 1076
- Verifying Router Advertisements | 1077

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 form 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: Sendbcast-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 | **939**
- Design 3: IPv6 Addressing with NDRA | **998**
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA | **1041**

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

IN THIS SECTION

- Requirements | **1078**
- Overview | **1079**
- Configuration | **1081**
- Verification | **1103**

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 | **1079**

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

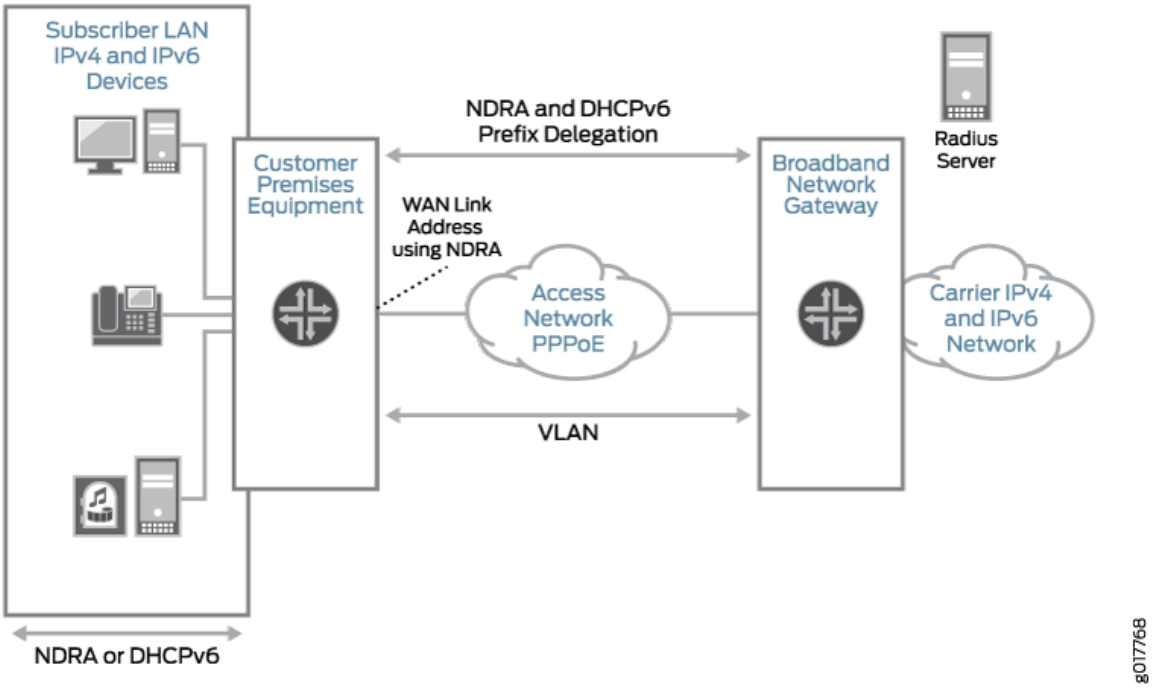

Table 61 on page 1080 describes the configuration components used in this example.

Table 61: 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.

Table 61: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation *(Continued)*

Configuration Component	Component Name	Purpose
	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 | **1082**
- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE | **1085**
- Configuring a Dynamic Profile for the PPPoE Logical Interface | **1086**
- Configuring a Loopback Interface | **1090**
- Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces | **1092**
- Specifying the BNG IP Address | **1094**
- Configuring RADIUS Server Access | **1095**
- Configuring RADIUS Server Access Profile | **1097**
- Specifying the RADIUS Server Access Profile to Use | **1099**
- Configuring Local Address-Assignment Pools | **1100**
- Specifying the Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | **1102**

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-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 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 | **1103**
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance | **1104**
- Verifying Dynamic Subscriber Sessions | **1105**
- Verifying DHCPv6 Address Pools Used for NDRA and DHCPv6 Prefix Delegation | **1106**
- Verifying DHCPv6 Address Bindings | **1108**
- Verifying Router Advertisements | **1108**
- Verifying the Status of the PPPoE Logical Interface | **1109**

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: Sendbcast-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 | **997**

DHCPv6 Prefix Delegation over PPPoE | **950**

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

IN THIS SECTION

- Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services | **1111**

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 IPv4, all traffic from single household has same source IPv4 public address. Hence, for every household only one IP demux IFL is created. For business subscribers, multiple public IP addresses are created using framed-routes. For IPv6, the source address of traffic coming from same household is different as each device has different source address. For scaling reason, it is not possible to have one IP demux IFL for each device in each household. Hence, all devices from the same household share the same IP demux IFL.

A TCP traffic received from the clients cannot trigger a subscriber session.

> **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.

Benefits of IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services

- 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 ALSO

> 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*.

NOTE: 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 log in from the new device with the same IP address. For example, a subscriber might log in from a laptop and then try to log in from a second laptop. 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 "Configuring Subscriber Session Timeout Options" on page 511.

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
```

SEE ALSO

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

RELATED DOCUMENTATION

Dual-Stack Access Models in a PPPoE Network | 1012

Dual Stack for PPPoE Access Networks Using NDRA | 1048

Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network | 1039

IPv6 WAN Link Addressing with NDRA | 939

Designs for IPv6 Addressing in a Subscriber Access Network | 993

Subscriber LAN Addressing with DHCPv6 Prefix Delegation | 948

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 | 1115
- On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview | 1115
- On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview | 1117
- IPCP Negotiation with Optional Peer IP Address | 1120
- How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled | 1121
- Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1122
- Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1122
- Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1123
- Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes | 1123
- Enabling IPv4 Release Control VSA (26–164) in RADIUS Messages | 1124

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 | **1115**
- IPv4 Address Release for Static PPP Subscribers | **1117**

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 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 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 | **1118**
- IPv4 Address Release for Dynamic PPP Subscribers | **1119**

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 pp0 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 | 1012
Dual Stack for PPPoE Access Networks Using DHCP | 1044
Dual Stack for PPPoE Access Networks Using NDRA | 1048

Dual Stack Subscribers Monitoring and Management

IN THIS SECTION

- Monitoring Active Subscriber Sessions | 1125
- Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance | 1126
- Monitoring Dynamic Subscriber Sessions | 1127
- Monitoring Address Pools Used for Subscribers | 1128
- Monitoring Specific Subscriber Sessions | 1130
- Monitoring the Status of the PPPoE Logical Interface | 1132
- Monitoring Service Sessions for Subscribers | 1133
- Monitoring PPP Options Negotiated with the Remote Peer | 1134
- Monitoring the RADIUS Attribute Used for NDRA | 1135

Monitoring Active Subscriber Sessions

IN THIS SECTION

- Purpose | 1125
- Action | 1126
- Meaning | 1126

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 | **1126**
- Action | **1127**
- Meaning | **1127**

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 | **1127**
- Action | **1127**
- Meaning | **1128**

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 | **1129**

Action | 1129
Meaning | 1130

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 | **1130**
- Action | **1131**
- Meaning | **1131**

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 | 1132
- Action | 1132
- Meaning | 1133

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: Sendbcast-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 | 1133
- Action | 1133
- Meaning | 1134

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 | 1134
- Action | 1134

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 | 1135
- Action | 1135

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 | **1000**

Dual-Stack Access Models in a PPPoE Network | **1012**

4

PART

Address-Assignment Pools for Subscriber Management

CHAPTER 8

Address-Assignment Pools for Subscriber Management

IN THIS CHAPTER

- Address-Assignment Pools for Subscriber Management | **1138**

Address-Assignment Pools for Subscriber Management

IN THIS SECTION

- Address-Assignment Pools Overview | **1139**
- Address Allocation from Linked Address Pools | **1141**
- Address-Assignment Pool Configuration Overview | **1146**
- Configuring an Address-Assignment Pool Name and Addresses | **1148**
- Configuring a Named Address Range for Dynamic Address Assignment | **1148**
- Preventing Addresses from Being Allocated from an Address Pool | **1149**
- Configuring Address-Assignment Pool Usage Threshold Traps | **1151**
- Configuring Address-Assignment Pool Linking | **1153**
- Configuring Address-Assignment Pool Hold-Down | **1154**
- Configuring DHCP Local Address Pool Rapid Drain | **1155**
- Configuring Static Address Assignment | **1157**
- Configuring Duplicate IPv4 Address Protection for AAA | **1158**
- Example: Configuring an Address-Assignment Pool | **1160**

Address-Assignment Pools Overview

IN THIS SECTION

- Address Assignment Types | 1139
- Named Address Ranges in Address Assignment Pool | 1139
- Address Allocation from Linked Address Pools | 1140
- Address Pool Hold-Down State | 1140
- Address-Assignment Pool for Neighbor Discovery Router Advertisement | 1140
- Excluding Specified Address or Address Range | 1141
- Licensing Requirement | 1141
- Benefits of Address Assignment Pools | 1141

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.

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..

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

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.

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 190.0.2.4 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 1143 shows the default search behavior.

Figure 26: Default Address Assignment from Linked Address Pools

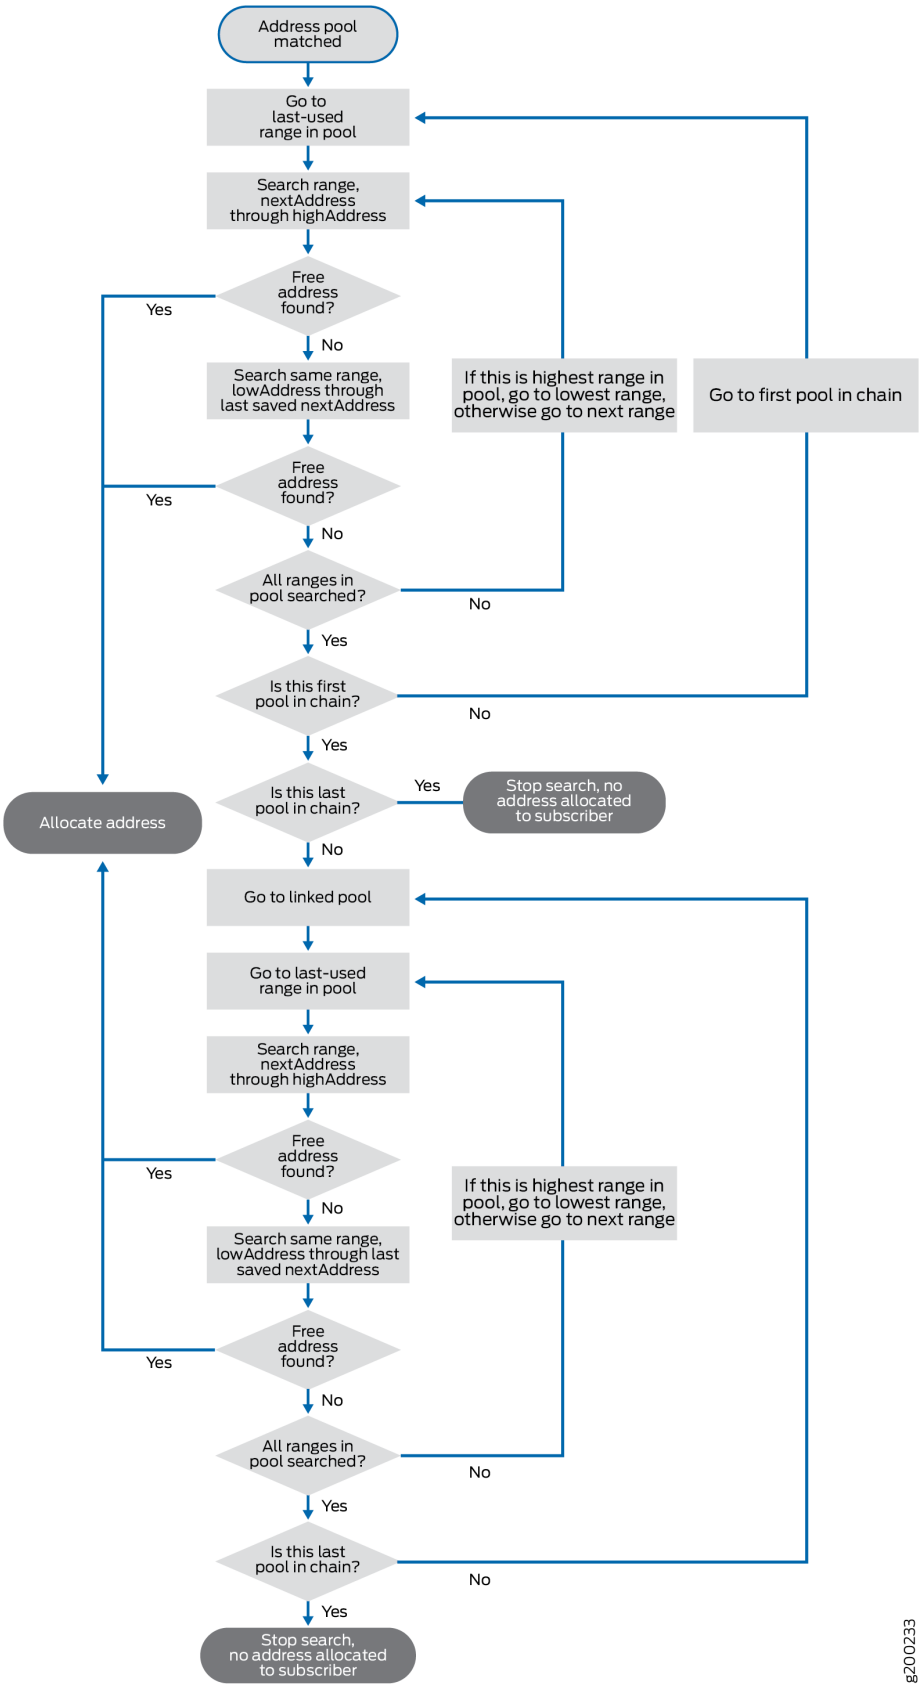

For example, suppose the following conditions exist:

- 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 1145 shows the search behavior when you include the `linked-pool-aggregation` statement.

Figure 27: Aggregated Address Assignment from Linked Address Pools

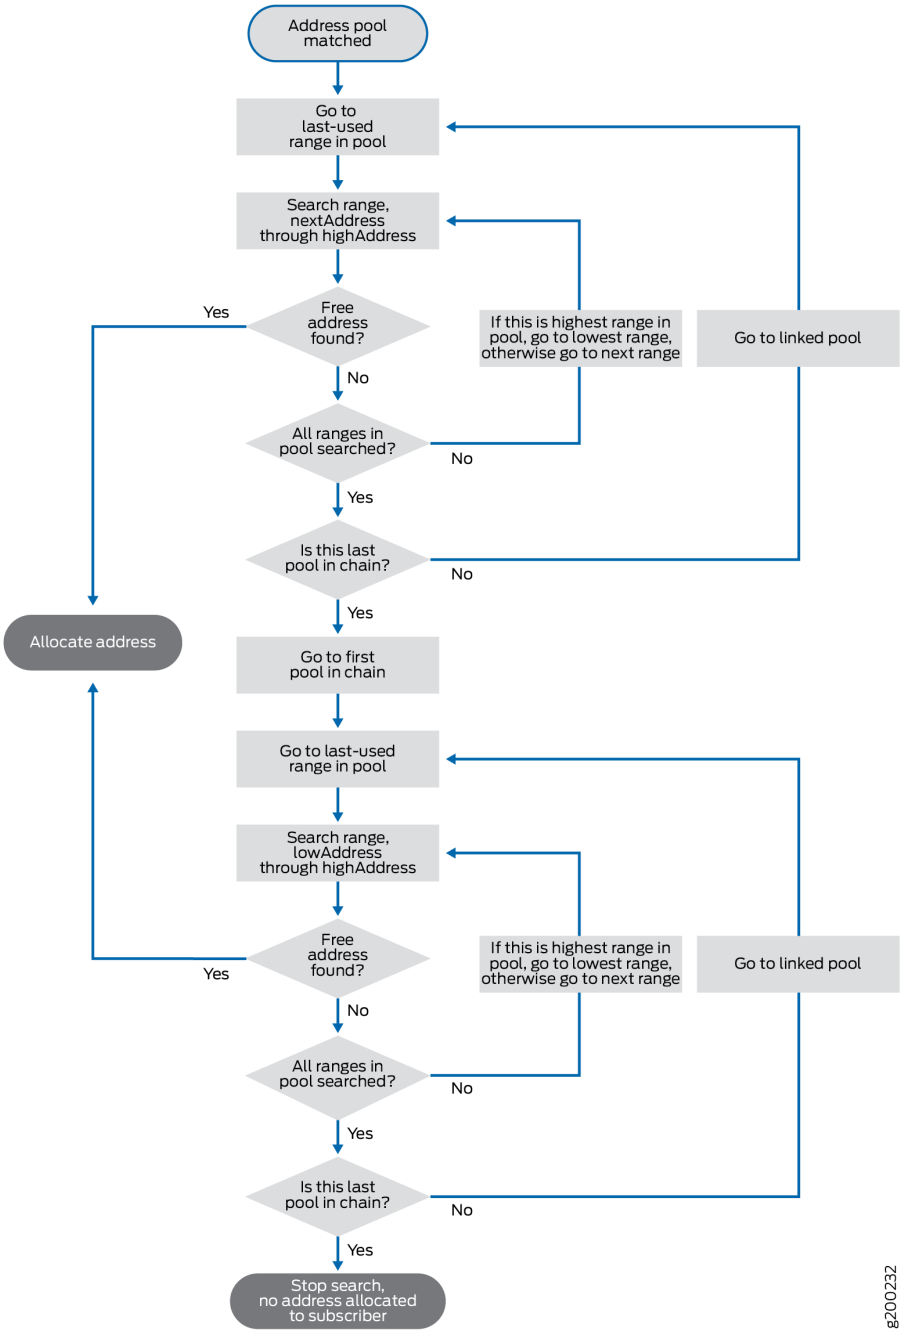

For example, consider the same conditions exist as for the default example:

- 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 1148.
2. (Optional) Configure named ranges (subsets) of addresses.

   See "Configuring a Named Address Range for Dynamic Address Assignment" on page 1148.
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 1149.
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 1153.
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 1154.
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 1155.
7. (Optional) Create static address bindings (IPv4 only).

   See "Configuring Static Address Assignment" on page 1157.
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 1158.
9. (Optional) Configure attributes for DHCP clients.

   See "Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address" on page 772.

SEE ALSO

DHCP Attributes Overview | 766

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 1146.

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` 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 | 772

Attributes That Can Be Applied to DHCP Clients | 768

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 svale6_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 | 1051

Example: Configuring an Address-Assignment Pool

IN THIS SECTION

- Requirements | 1161
- Overview | 1161
- Configuration | 1161

Requirements

Overview

Configuration

IN THIS SECTION

- CLI Quick Configuration | **1161**

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`.

Release History Table

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 `reassign-on-match` 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.

PART 5

DNS Addresses for Subscriber Management

DNS Addresses for Subscriber Management | 1165

CHAPTER 9

DNS Addresses for Subscriber Management

IN THIS CHAPTER

- DNS Name Server Addresses for Subscriber Management | 1165

DNS Name Server Addresses for Subscriber Management

IN THIS SECTION

- DNS Name Server Address Overview | 1165
- Configuring DNS Name Server Addresses for Subscriber Management | 1167
- Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | 1169
- DNS Resolver for IPv6 DNS Overview | 1170
- Configuring a DNS Server Address for IPv6 Hosts | 1170

DNS Name Server Address Overview

IN THIS SECTION

- Benefits of Local DNS Server Addresses | 1166

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 specify 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 | 768

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 772

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 1138. 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 1165.

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 | 768

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address | 772

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 | 957

DHCPv6 Options in a DHCPv6 Multiple Address Environment | 958

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.

PART 6

M:N Subscriber Redundancy

M:N Subscriber Redundancy | 1173

CHAPTER 10

M:N Subscriber Redundancy

IN THIS CHAPTER

- M:N Subscriber Redundancy on BNG | **1173**
- M:N Subscriber Service Redundancy on DHCP Server | **1219**
- N+1 Support for BNG M:N Subscriber Service Redundancy | **1223**
- BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | **1227**

M:N Subscriber Redundancy on BNG

IN THIS SECTION

- M:N Subscriber Redundancy on BNG Overview | **1173**
- How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization | **1205**
- How to Configure M:N Subscriber Redundancy with Pseudowires and DHCP Binding Synchronization | **1211**
- Verifying M:N Redundancy and Active Leasequery Topology Discovery Information | **1216**

M:N Subscriber Redundancy on BNG Overview

IN THIS SECTION

- Benefits of M:N Subscriber Redundancy on BNG | **1176**
- Fundamentals of M:N Redundancy | **1177**
- Subscriber Sessions and Hot Standby Mode | **1180**
- M:N Redundancy Using Virtual Router Redundancy Protocol (VRRP) | **1181**

- VRRP Failover and Reversion Timing | **1184**
- M:N Redundancy Using Pseudowire Redundancy | **1184**
- DHCP Active Leasequery Topology Discovery and M:N Subscriber Redundancy | **1186**
- Example Topology Discovery with VRRP Redundancy | **1189**
- Example Topology Discovery with Pseudowire Redundancy | **1192**
- Static Subscribers and M:N Redundancy | **1196**
- Convergence and M:N Subscriber Redundancy | **1202**

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 1175.

Figure 28: Sample Topology for M:N Subscriber Group Redundancy

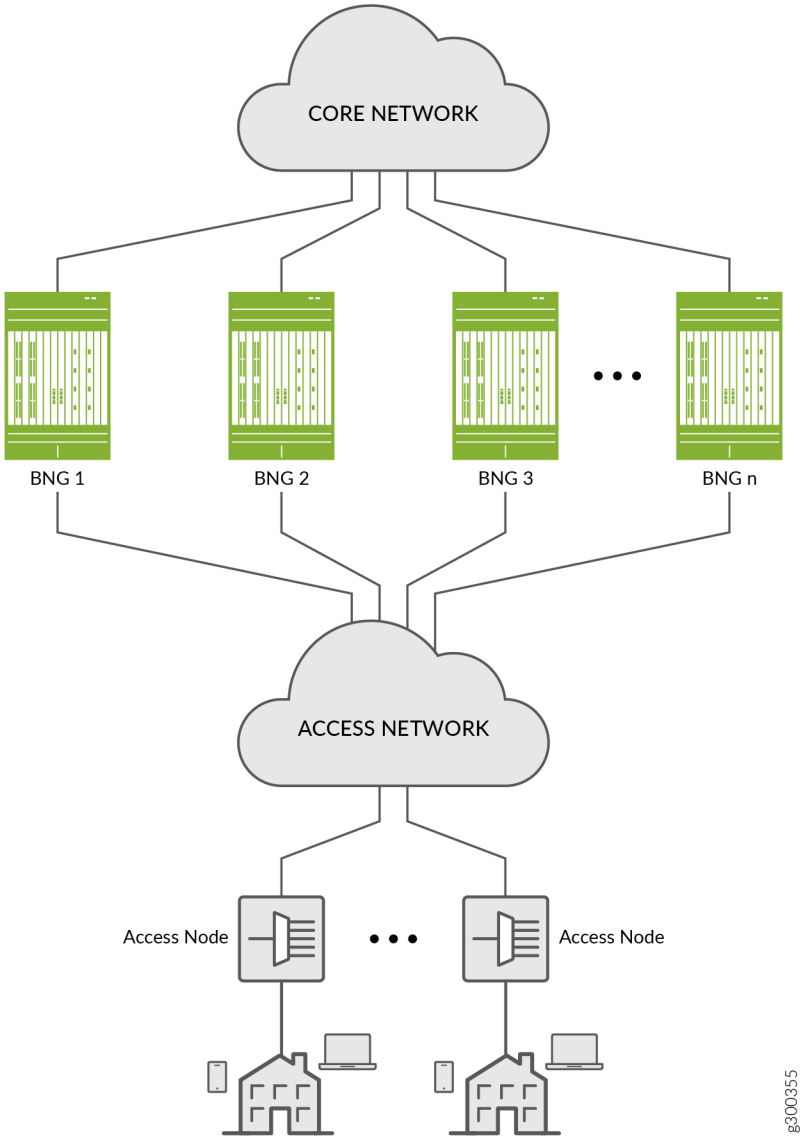

A failure in any of the locations listed in Table 62 on page 1175 can trigger a primary BNG to fail over to a backup BNG.

Table 62: Types of Failures Mitigated by M:N Subscriber Group Redundancy

Access line card	Core-facing link
Access link	Partial access network

Chassis	Partial core 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 1175 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 1178 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

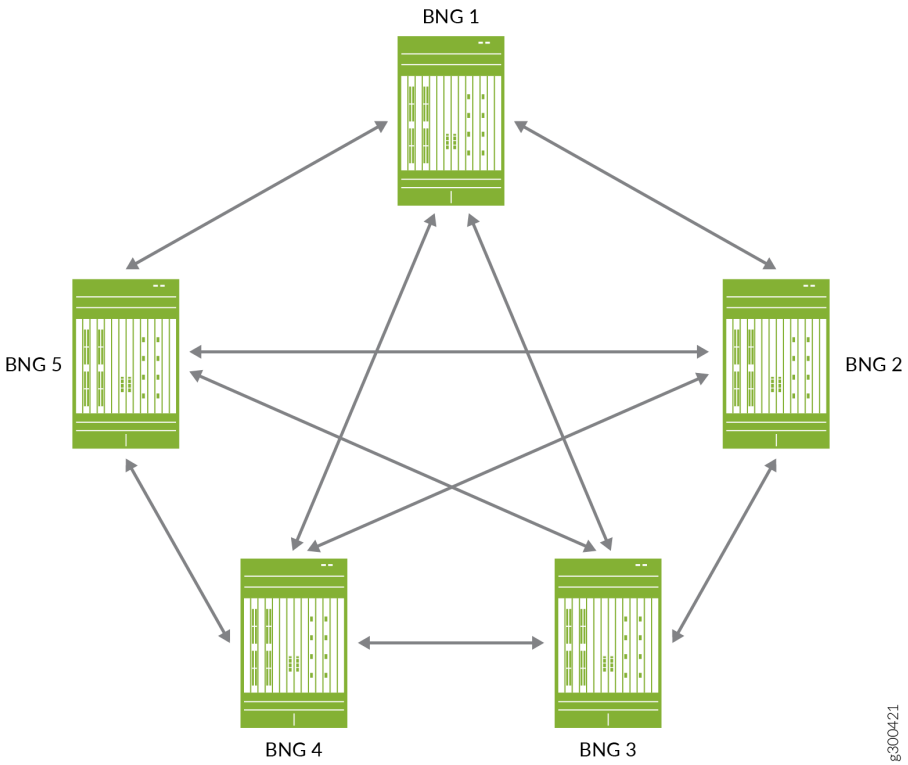

For M:N redundancy, it is important to understand that you can configure:

- 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 primarys for the groups.

Figure 30 on page 1179 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

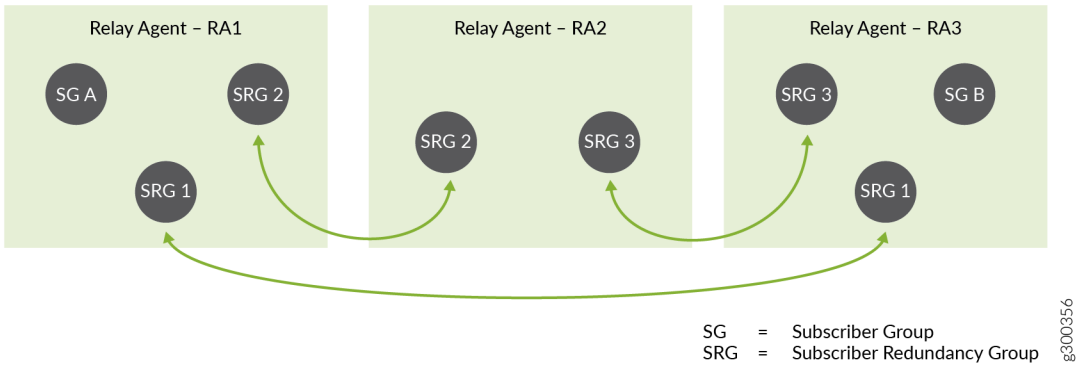

- Relay agent RA1 is configured for subscriber redundancy groups, SRG 1 and SRG 2, and subscriber group SG A.
- Relay agent RA2 is configured for SRG 2 and SRG 3.
- Relay agent RA3 is configured for SRG 1, SRG 3, and SG B.

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 1180, 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

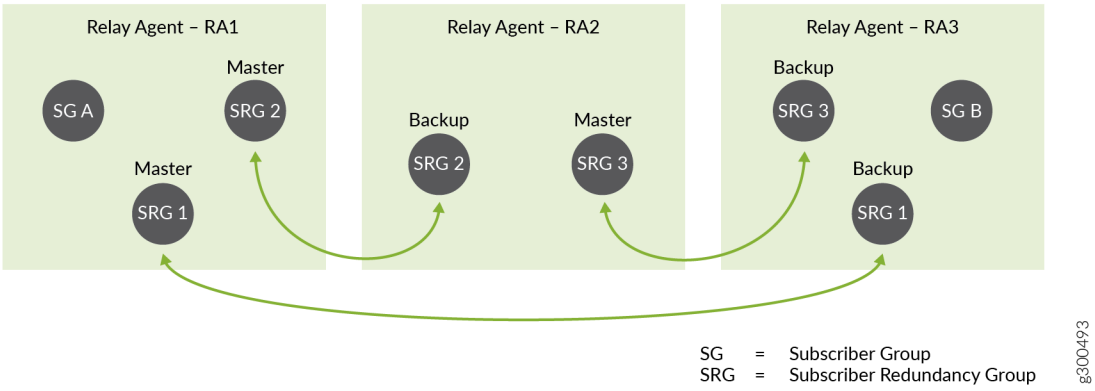

If this BNG fails, then it fails over to a different backup BNG for SRG 1 and SRG 2, as shown in Figure 32 on page 1180.

Figure 32: Primary and Backup BNGs for Subscriber Redundancy Groups After Failover

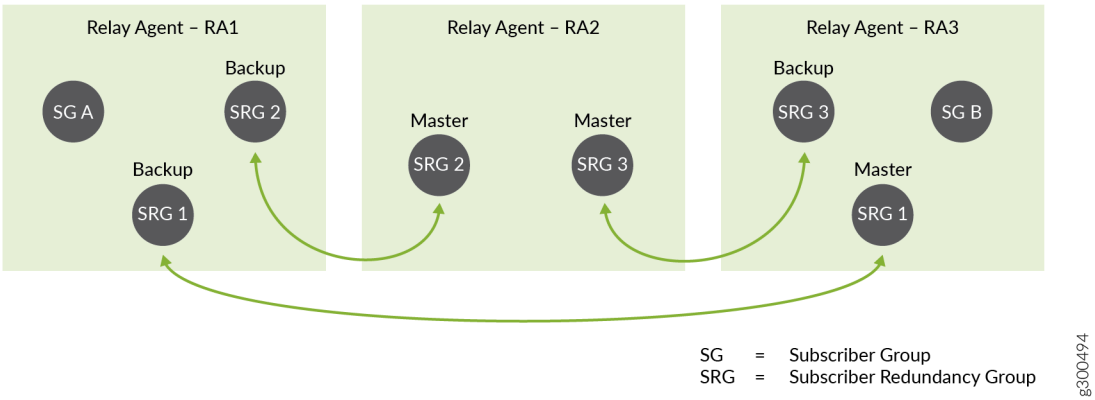

- For SRG 1, it fails over to the BNG hosting RA 3. The RA 3 BNG becomes the new primary for SRG 1.
- For SRG 2, it fails over to the BNG hosting RA 2. The RA 2 BNG becomes the new primary for SRG 2.

The failure has no effect on SRG 3.

Subscriber Sessions and Hot Standby Mode

Each backup BNG is in hot-standby mode for it's 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.

Figure 33 on page 1183 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

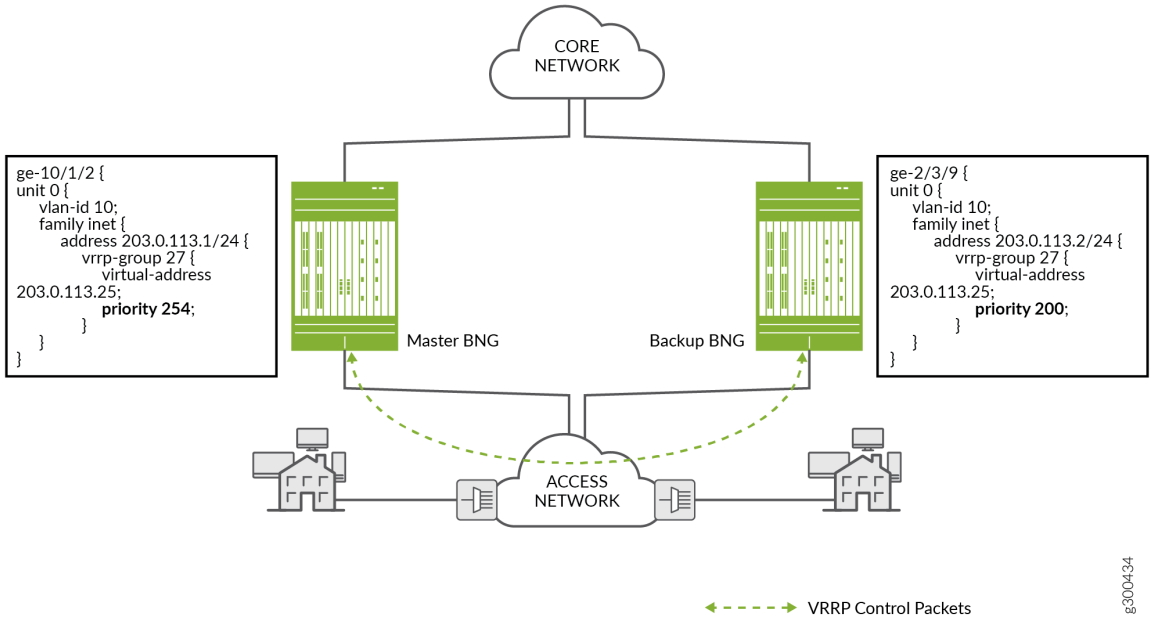

Figure 34 on page 1183 shows how the configured VRRP priority determines which BNG acts as the primary or backup for a subscriber redundancy group.

Figure 34: VRRP Priorities for Three Subscriber Redundancy Groups

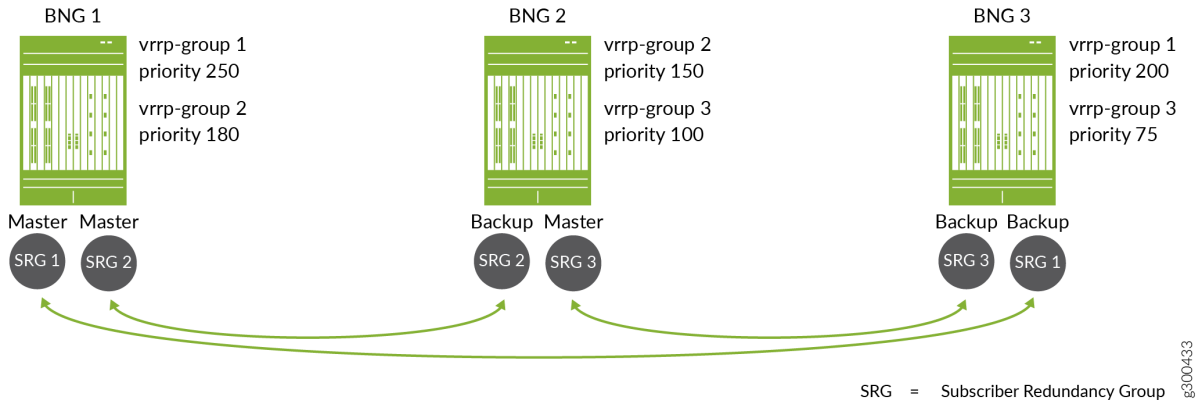

The topology includes three subscriber redundancy groups (*M*), SRG 1, SRG 2, and SRG 3 on three BNGs (*N*). Each subscriber redundancy group corresponds to a different VRRP group. The arrows indicate the primary router and backup router for each group

- 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 1183, 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.

Figure 35 on page 1185 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

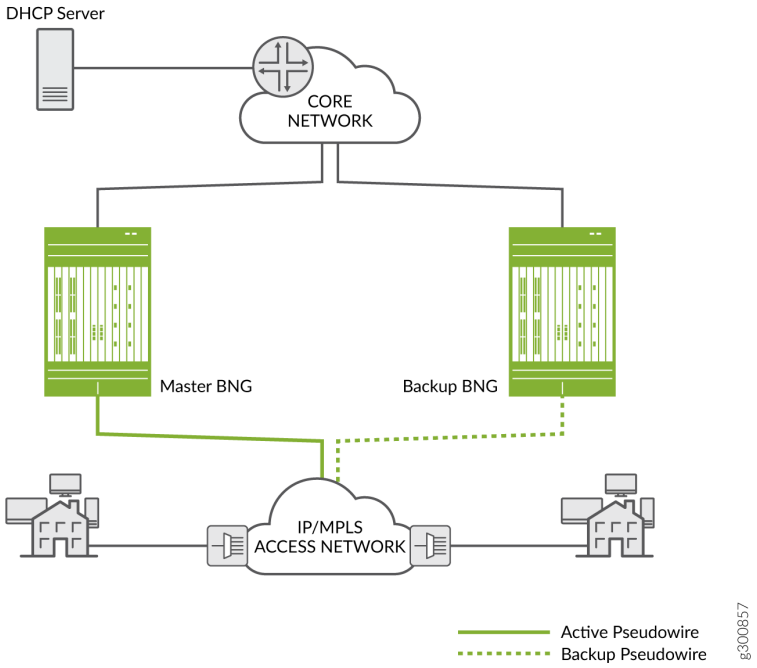

For L2 circuits, you configure the pseudowires as the underlying (access-facing) interfaces on the BNGs. You then configure the interfaces with L2 connections such as Ethernet, dynamic auto-sensed VLANs, or static VLANs. The DHCP client-facing, pseudowire interfaces are bundled and added to a L2 circuit (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 1188 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

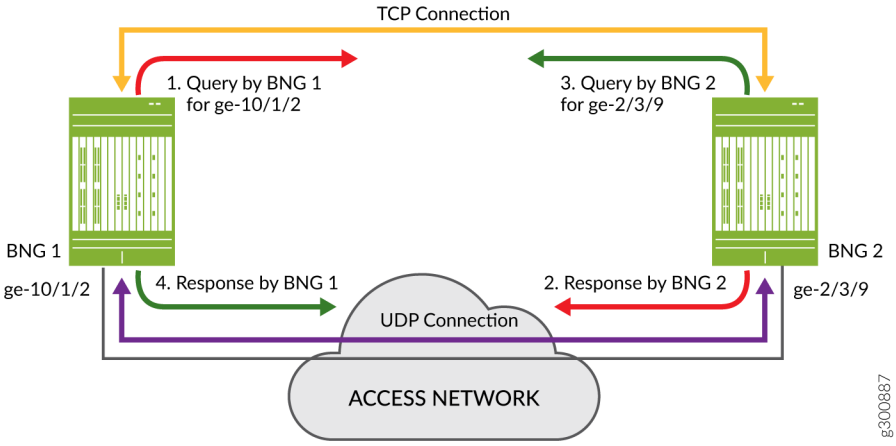

Figure 37 on page 1188 shows a query and response for two DHCP relay agents on BNGs when you use pseudowire redundancy. R1 sends the query for its access interface, ps2.0, to BNG 2 over the TCP connection. R2 responds over the same TCP connection. R2 also sends a query to R1, for its access interface, 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

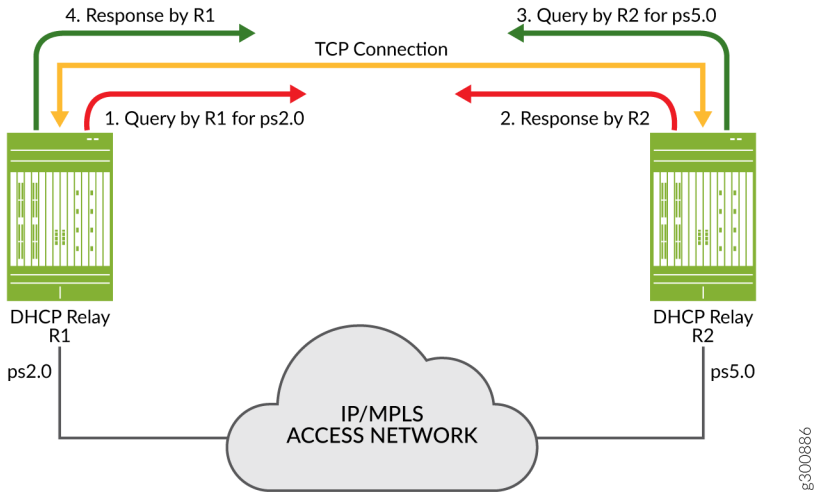

Each peer agent sends queries to its peers so it can build its own translation table of corresponding local and remote access interfaces. In this way all relay agents that you configure both as peers and for

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 1190 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

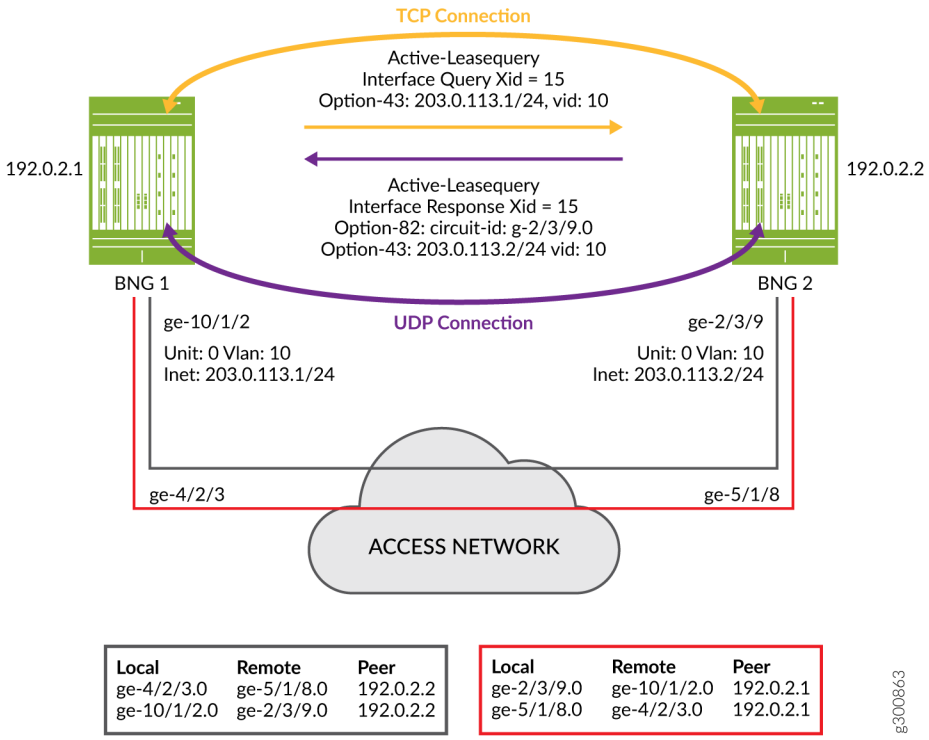

1. After TCP synchronization, peer 192.0.2.1 sends a topology discovery query to peer 192.0.2.2 to determine the matching remote interface for its own local interface, ge-10/1/2.0. Because this is a DHCPv4 topology, the message it sends is a DHCPLEASEQUERY. The query is sent over the TCP connection and 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.

2. 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 1190 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 799 and "Configuring and Using DHCP Active Leasequery" on page 818. The *Topology Discovery Messages* section in "DHCP Active Leasequery" on page 799 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 1193 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

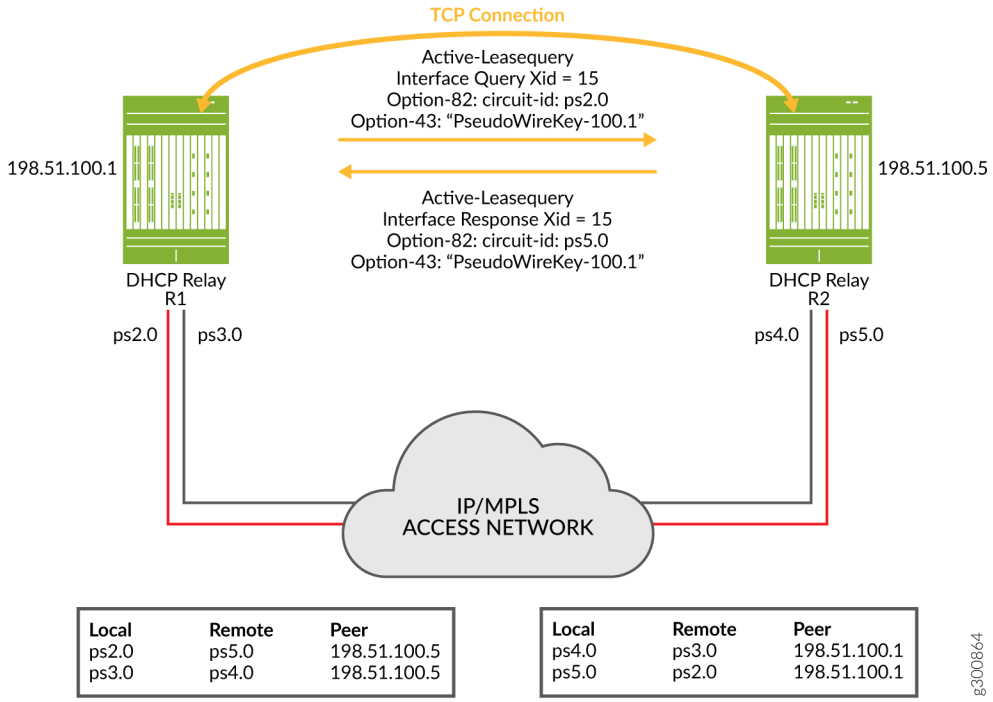

1. After TCP synchronization, peer 198.51.100.1 sends a topology discovery query to peer 198.51.100.5 to determine the matching remote interface for its own local interface, ps2.0. Because this is a DHCPv4 topology, the message it sends is a DHCPLEASEQUERY. The query is sent over the TCP connection and 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 1193. 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 1195. 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

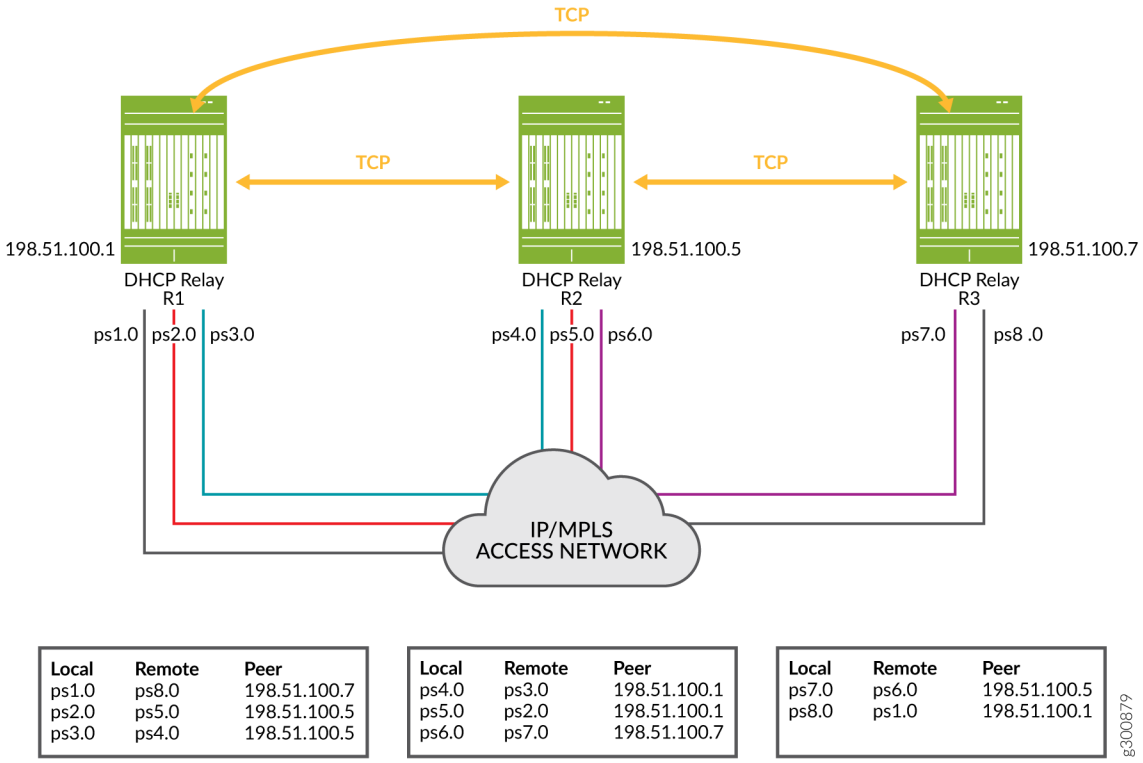

Compare the translation tables and colored pseudowire connection lines in Figure 40 on page 1195 with the shared key configuration snippets for each relay agent in Figure 41 on page 1195.

Figure 41: Sample Shared Key Configuration for Three Peers

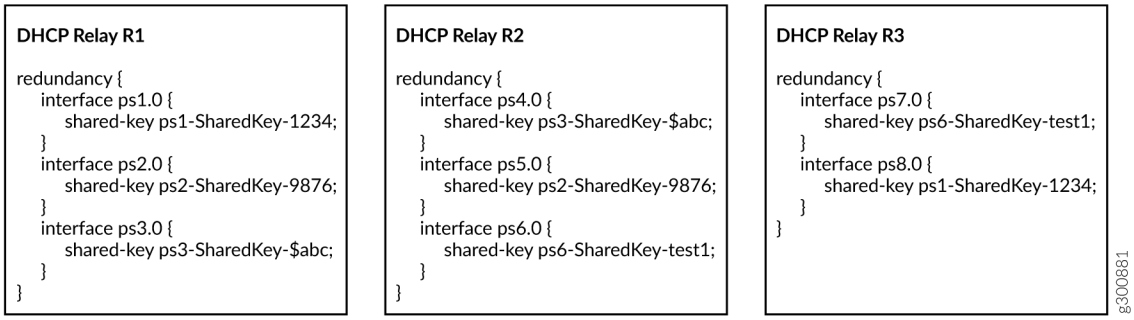

You can see that interface ps1.0 on R1 has the same shared key as interface ps8.0 on R3. The translation tables for R1 and R3 show this relationship was discovered by the topology discovery process.

Similarly, interface ps2.0 on R1 and ps5.0 on R2 have the same shared key. Again, topology discovery determined this relation ship 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 799 and "Configuring and Using DHCP Active Leasequery" on page 818. The *Topology Discovery Messages* section in "DHCP Active Leasequery" on page 799 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 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 theprimary 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 | **1206**
- Configure VRRP to Support M:N Redundancy | **1208**
- Configure Active Leasequery with Topology Discovery | **1210**

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 788
- 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 814.

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" on page 2278 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 | **1212**
- Configure Active Leasequery with Topology Discovery | **1215**

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 788
- *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 814.

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" on page 2278 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 | 1217
- Action | 1217

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)
```

Release History Table

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 | **788**
- DHCP Active Leasequery | **799**

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 | **1219**

M:N Subscriber Service Redundancy on DHCP Server Overview

IN THIS SECTION

- Benefits of M:N Subscriber Service Redundancy on DHCP Server | **1223**

You can configure M:N subscriber service redundancy on DHCP server running on MX Serirs 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 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 1220.

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 1221 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

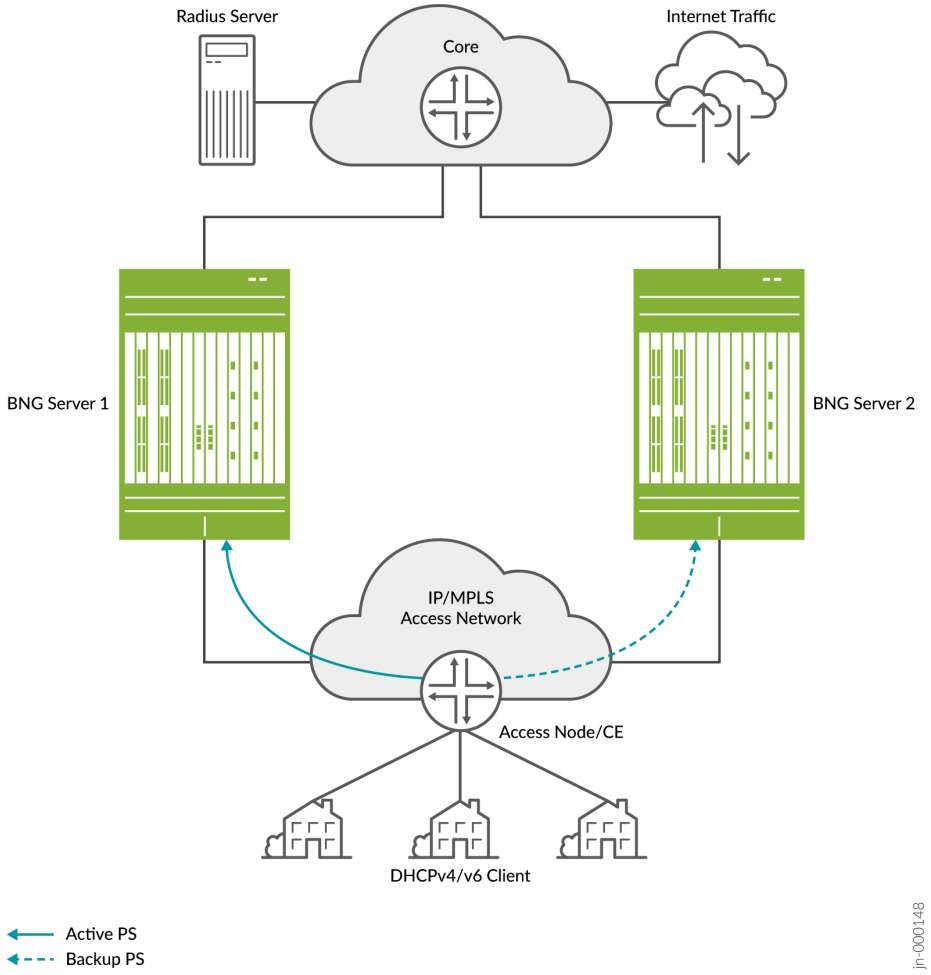

Figure 43 on page 1222 shows the topology for L2 circuit based IP/MPLS PWHT in client-relay-server mode.

Figure 43: L2 Circuit Based IP/MPLS PWHT in Client-Relay-Server Mode

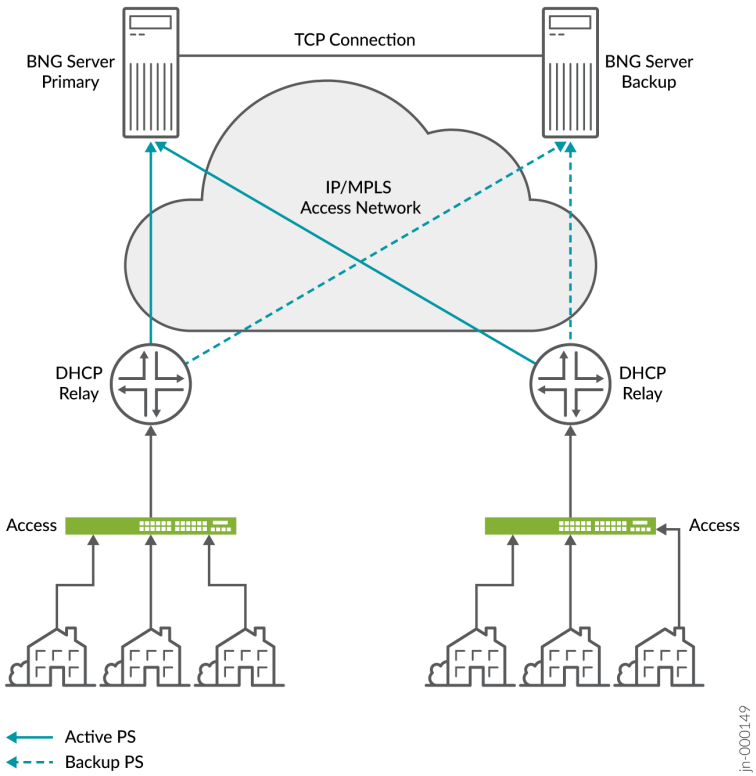

In both client-server mode and client-relay-server mode topology, BNG servers use TCP connection for active leasequery to synchronize binding details. The subscriber service redundancy on DHCP server occurs in the following order:

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 folows:

- 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

active-leasequery (DHCP Local Server) | **1577**

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 | **1224**
- How N+1 Support for BNG M:N Subscriber Service Redundancy Works | **1224**

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

IN THIS SECTION

- Benefits of N+1 Support for BNG M:N Subscriber Service Redundancy | **1224**

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. 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.

Benefits of N+1 Support for BNG M:N Subscriber Service Redundancy

- Reduces the cost of deploying backup BNGs.

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

IN THIS SECTION

- Subscriber Service Redundancy When Primary BNG Fails | **1226**
- Subscriber Service Revert When the Primary BNG Becomes Active | **1226**

Figure 44 on page 1225 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

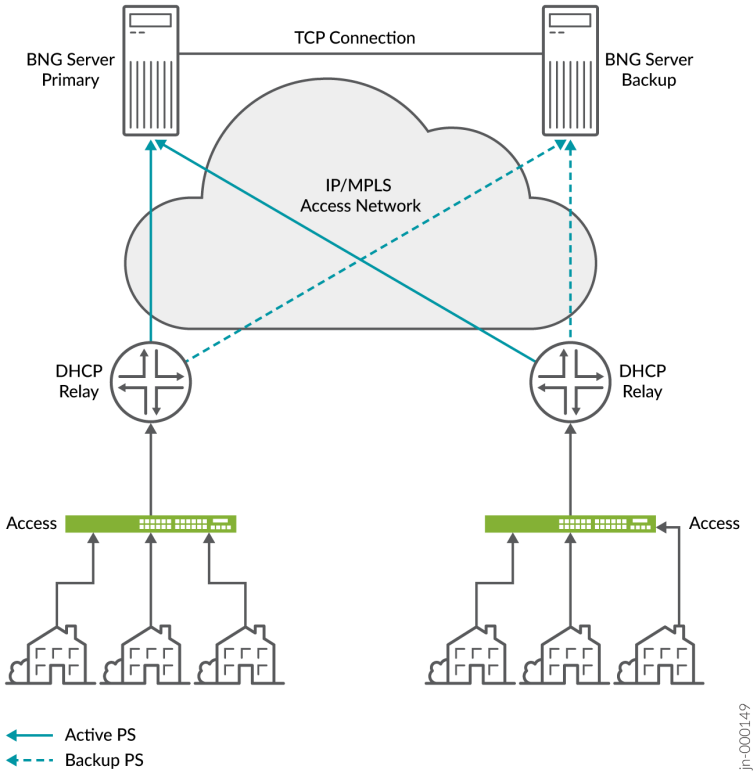

When the subscribers log in to the primary BNGs, the active leasequery brings the subscriber state to the backup BNG. As the backup BNG hosts the `service-activation-on-failover` mode, the backup BNG consumes minimal Packet Forwarding Engine resources and back ups up to 192000 subscribers.

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 1225, 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 | 3124

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 | **1227**
- How BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery Works | **1228**

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

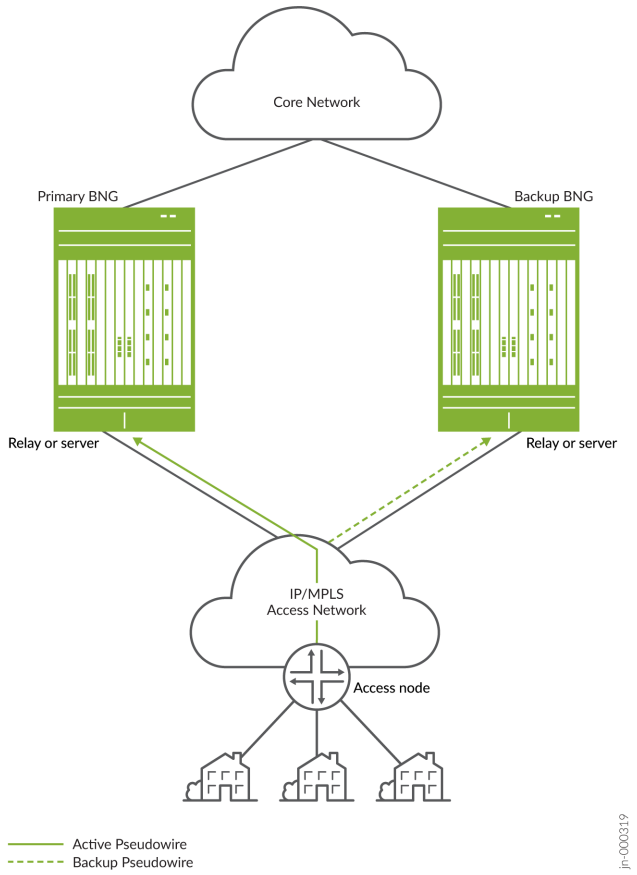

Figure 45 on page 1229 shows the topology diagram for layer 2 circuit based on IP/MPLS pseudowire headend termination (PWHT) scenario.

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).

RELATED DOCUMENTATION

auto-configure (IPv4) | **1647**

auto-configure (IPv6) | **1650**

PART 7

Access Node Control Protocol and the ANCP Agent for Subscriber Services

CHAPTER 11

Access Node Control Protocol and the ANCP Agent for Subscriber Services

IN THIS CHAPTER

- ANCP Agent Neighbors and Operations | **1232**
- ANCP Agent Traffic Shaping and CoS | **1292**
- ANCP Agent and AAA | **1311**
- ANCP Monitoring and Management | **1326**
- Tracing ANCP Events for Troubleshooting | **1333**

ANCP Agent Neighbors and Operations

IN THIS SECTION

- ANCP and the ANCP Agent Overview | **1233**
- ANCP Operations in Different Network Configurations | **1243**
- Configuring the ANCP Agent | **1253**
- Configuring ANCP Neighbors | **1255**
- Associating an Access Node with Subscribers for ANCP Agent Operations | **1256**
- Specifying the Interval Between ANCP Adjacency Messages | **1257**
- Specifying the Maximum Number of Discovery Table Entries | **1258**
- Configuring the ANCP Agent for Backward Compatibility | **1258**
- Specifying How Long Processes Wait for the ANCP Agent Restart to Complete | **1259**
- Configuring the ANCP Agent to Learn ANCP Partition IDs | **1260**
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet | **1261**

ANCP and the ANCP Agent Overview

IN THIS SECTION

- Overview | **1233**
- Topology Discovery | **1234**
- Subscriber Services | **1234**
- ANCP Interfaces and Access Loop Circuit Identifiers | **1235**
- Mapping Access Lines to Interfaces and Interface Sets | **1236**
- ANCP Neighbors | **1237**
- Partitions | **1239**
- Adjacency Update Messages | **1240**
- Generic Response Messages and Result Codes | **1240**
- Benefits of Access Node Control Protocol | **1242**

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" on page 1540 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" on page 2541
- "clear ancp subscriber" on page 2547

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 `gsmp-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 1241 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 | **1244**
- Business Services Traffic Shaping Model | **1246**
- ANCP Network Using N:1 and 1:1 Configuration Models without Interface Sets | **1247**
- Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets | **1248**
- ANCP Network Using N:1 Configuration Model with Interface Sets | **1249**
- Sequence of ANCP Events: Static VLAN Interfaces over Ethernet with Interface Sets | **1251**
- ANCP Network Using 1:1 Configuration Model with Interface Sets | **1251**
- Sequence of ANCP Events: Static VLAN Demux Interfaces over Aggregated Ethernet with Interface Sets | **1253**

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 1244 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 1245 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	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. The ANCP agent automatically maps the ACI from the discovery packet to the dynamically created interface set. All DHCP IP demux or PPPoE sessions that have the same ACI are mapped to the same interface set.	Dynamic VLAN and VLAN demux interfaces.	Required.
Dynamic interface sets	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. The ANCP agent configuration must include the `access-identifier` statement to statically map the ACI to the interface set. All DHCP IP demux and PPPoE sessions are mapped to an interface set according to the rules of the dynamic profile.	DHCP IP demux subscriber interfaces, PPPoE subscriber interfaces, or VLAN interfaces.	Irrelevant.
Static interface sets	The interface set and set name are statically configured and include multiple static interfaces. The ANCP agent configuration must include the `access-identifier` statement to statically map the ACI to the interface set.	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	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. The ANCP agent configuration must include the `access-identifier` statement to statically map the ACI to the interface set. All DHCP IP demux or PPPoE sessions that have the same VLAN ID tag are mapped to the same interface set.	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 dervices 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 46 on page 1248). 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 46: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model)

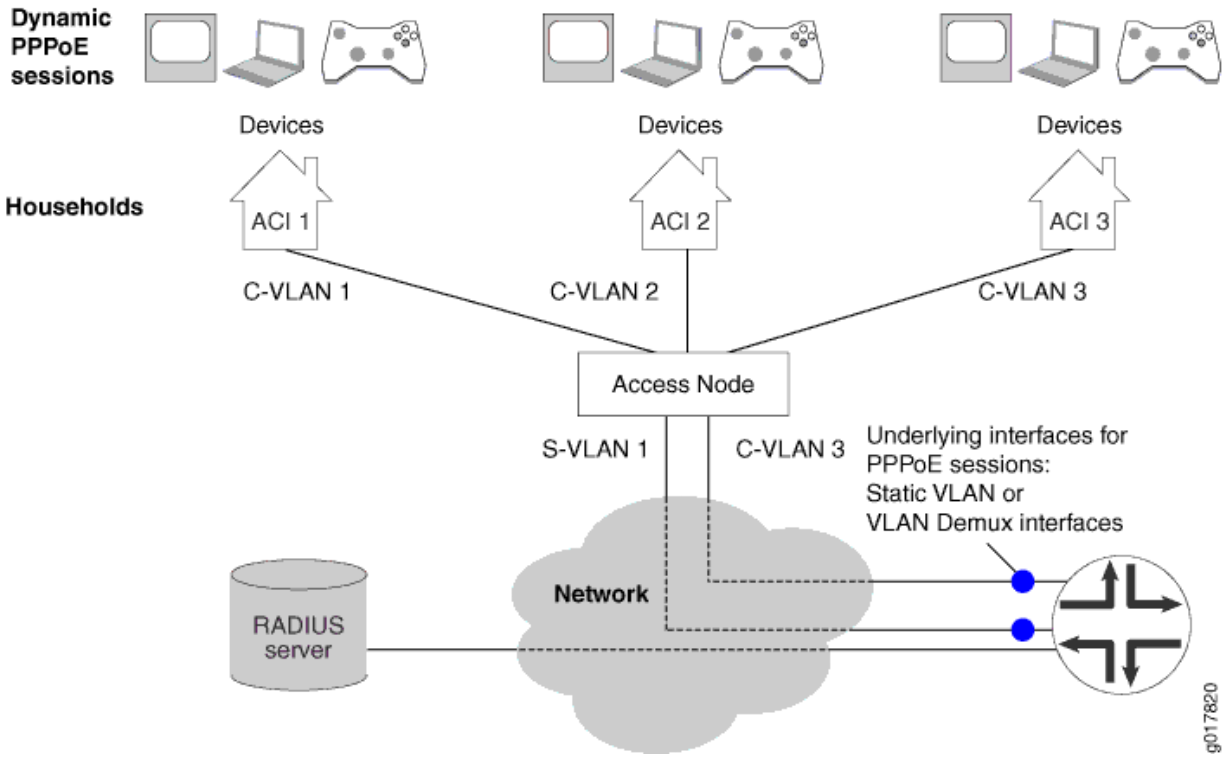

Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets

The following sequence of events is for the topology in Figure 46 on page 1248 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 47 on page 1250). 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 47: Sample ANCP Topology with Interface Sets (N:1 Model)

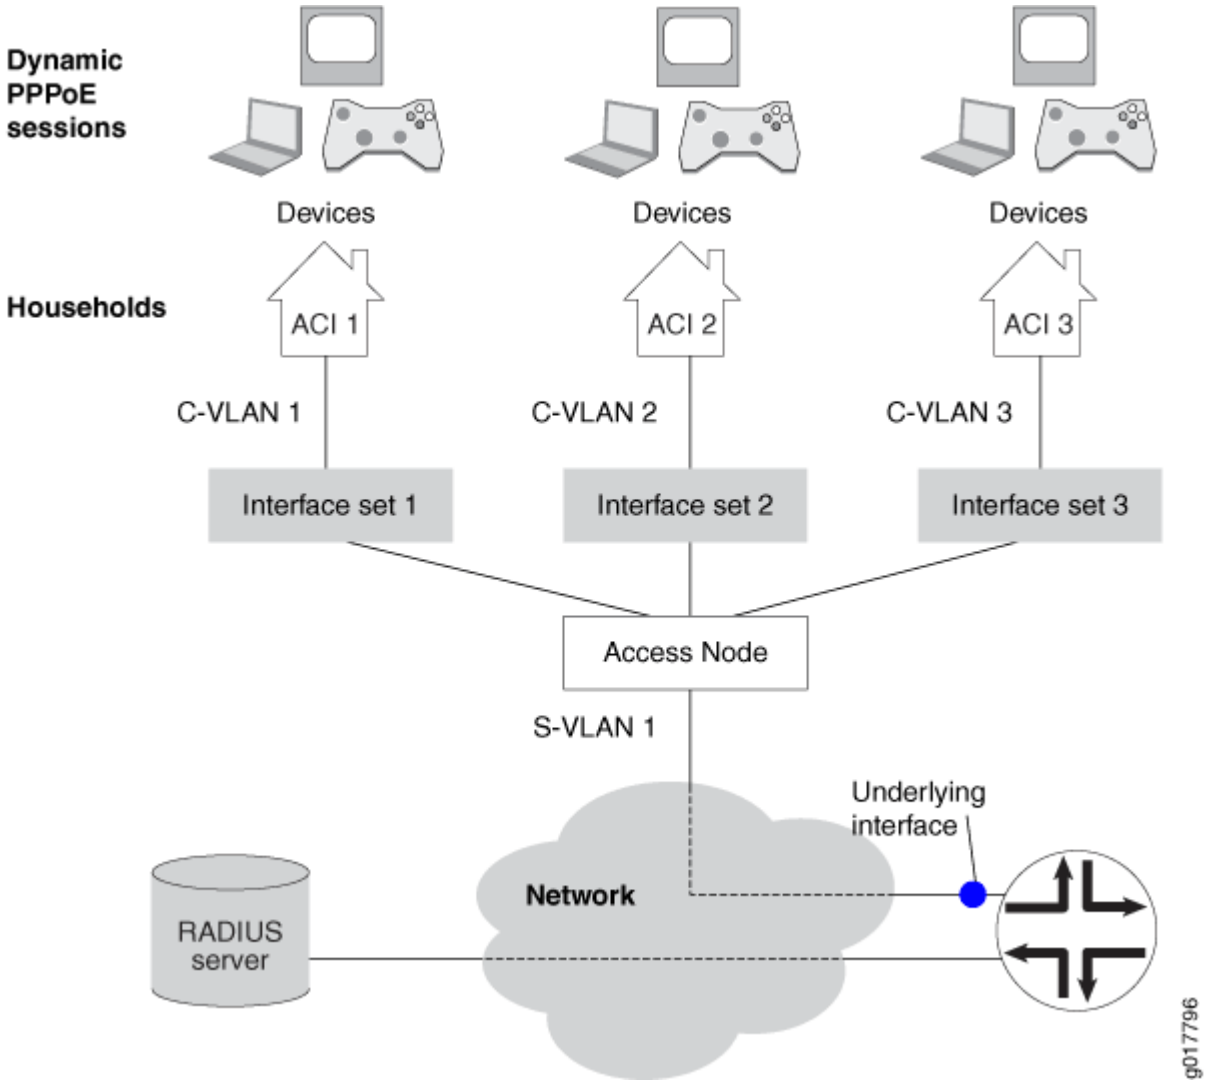

In this N:1 model with interface sets, the access node must add the DSL Forum VSA to the PPPoE PADI and PADR discovery packets that it passes to the router during the establishment of dynamic PPPoE sessions. The VSA includes the ACI for the household. This inclusion enables AAA to correlate the PPPoE 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 47 on page 1250 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 48 on page 1252). 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 48: Sample ANCP Topology with Interface Sets (1:1 Model)

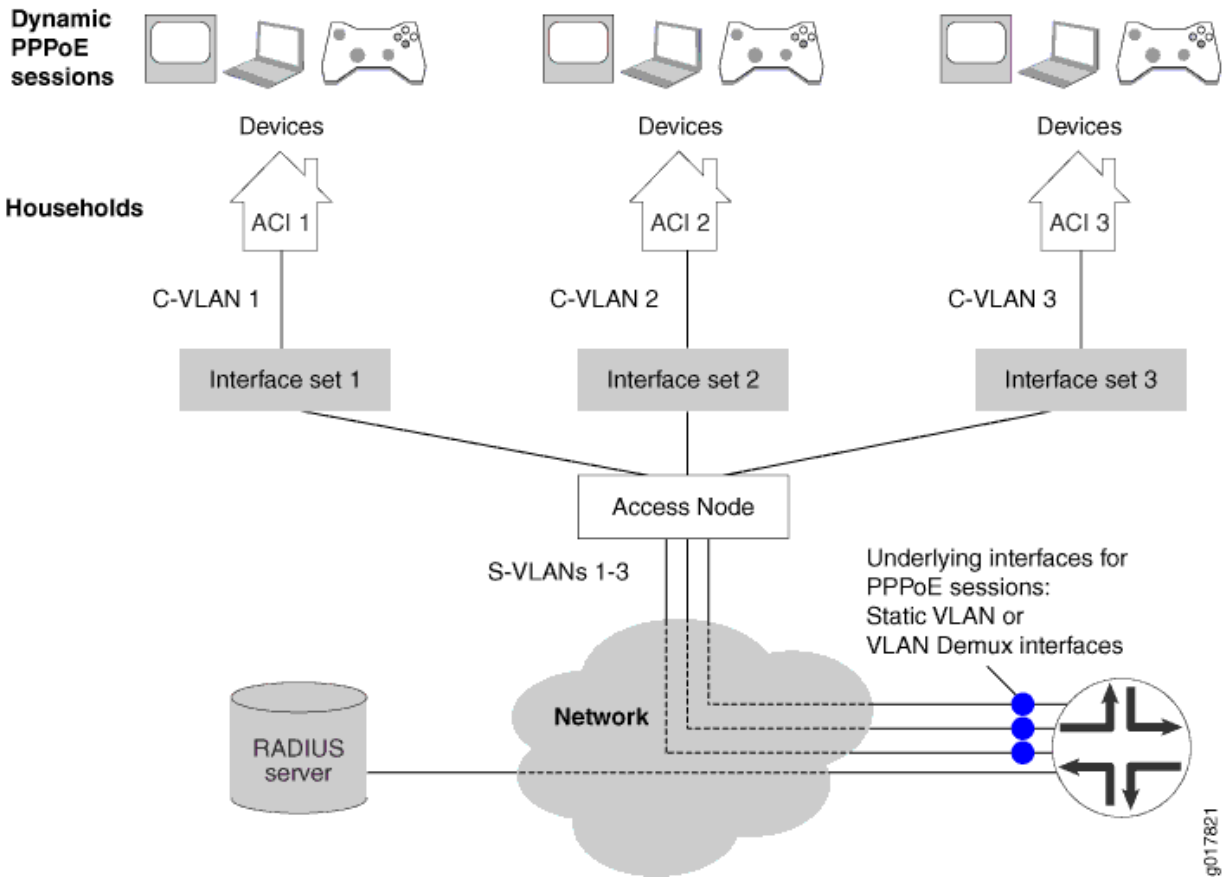

In this 1:1 model with interface sets, the ANCP agent configuration must map the underlying interface for the PPPoE sessions in an interface set to both the ACI and the interface set. This configuration enables AAA to correlate the PPPoE sessions with their respective subscriber access lines and DSL 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 48 on page 1252 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 1255.

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 1256.

3. (Optional) Configure the adjacency timer.

   See "Specifying the Interval Between ANCP Adjacency Messages" on page 1257.

4. (Optional) Specify the maximum number of discovery table entries that are accepted.

   See "Specifying the Maximum Number of Discovery Table Entries" on page 1258

5. (Optional) Configure the ANCP agent to work with an early IETF draft.

   See "Configuring the ANCP Agent for Backward Compatibility" on page 1258.

6. (Optional) Configure the graceful restart timer.

   See "Specifying How Long Processes Wait for the ANCP Agent Restart to Complete" on page 1259.

7. (Optional) Configure the ANCP agent to learn partition IDs from neighbors.

   See "Configuring the ANCP Agent to Learn ANCP Partition IDs" on page 1260.

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 1306.

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 1308.

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 1299.

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 1304.

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 1324.

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 1325.

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 1333.

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 interoperation 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 | 1261
- Overview | 1262
- Configuration | 1269
- Verification | 1287

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 1233
  - "ANCP Operations in Different Network Configurations" on page 1243
- Configured your access node.
- Configured your RADIUS server.

Overview

IN THIS SECTION

- Topology | 1263

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 49 on page 1263 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 49: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet

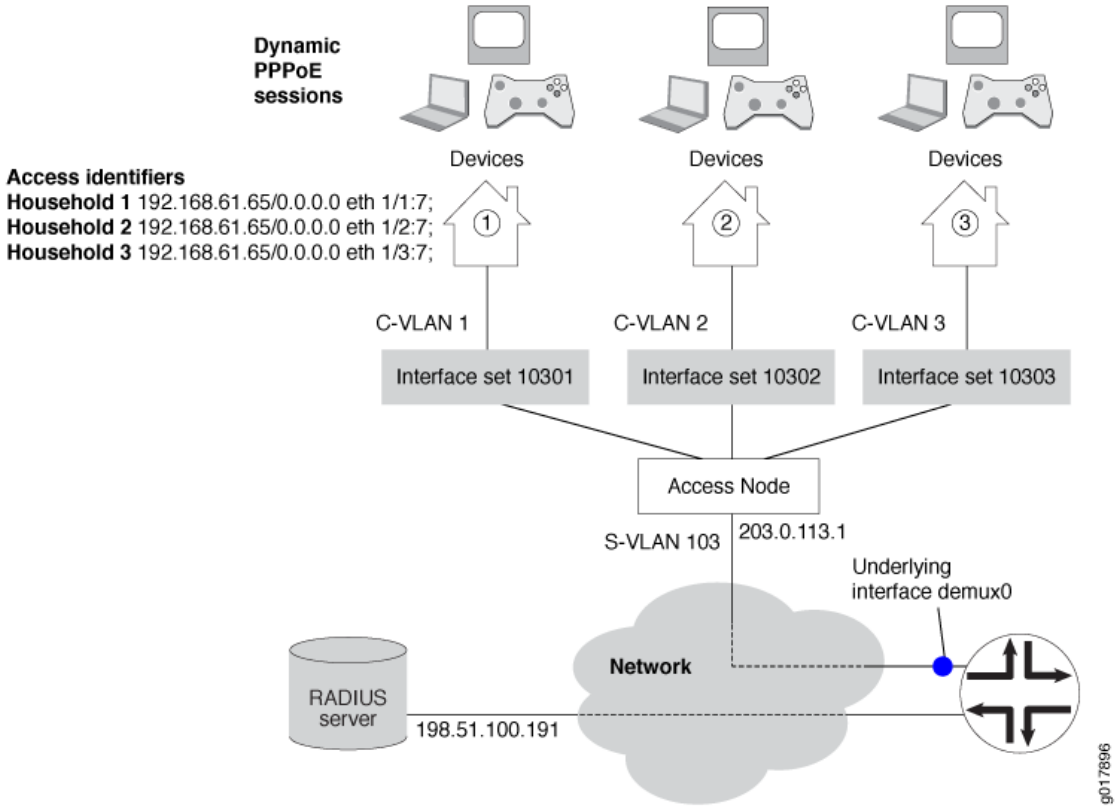

Table 67 on page 1264 describes the configuration components used in this example.

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 ancp-10302 ancp-10303	Each profile defines the dynamic PPPoE session created when any of the devices for a particular subscriber household accesses the network. Each profile specifies the following: • 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.
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	Aggregated Ethernet interface that is the underlying interface for the VLAN demux interfaces. The interface includes the following configuration: • 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	VLAN demux logical interfaces that correspond to the C-VLANs for individual subscriber households. Each logical interface includes the following configuration: • 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 | **1269**
- Configuring the Dynamic PPPoE Profiles | **1272**
- Configuring the Static VLAN Demux Interface over Aggregated Ethernet | **1275**
- Configuring Class of Service | **1281**
- Configuring ANCP | **1283**
- Configuring RADIUS Authentication and Accounting | **1285**

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;
```

```
                }
                pppoe-options {
                    underlying-interface "$junos-underlying-interface";
                }
                keepalives interval 30;
                family inet {
                    unnumbered-address lo0.0;
                }
            }
        }
    }
}
ancp-10302 {
    interfaces {
        interface-set 10302 {
            interface pp0 {
                unit "$junos-interface-unit";
            }
        }
        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;
                }
            }
        }
    }
}
ancp-10303 {
    interfaces {
        interface-set 10303 {
            interface pp0 {
                unit "$junos-interface-unit";
            }
        }
```

```
        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 gigether-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;
    gigether-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-adsl 90
user@host1# set qos-adjust-adsl2 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 | **1287**
- Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set | **1288**
- Verifying the demux0 Interface Configuration | **1289**
- Verifying the pp0 Interface Configuration | **1290**
- Verifying the ANCP Agent Configuration | **1291**

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: Sendbcast-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 | **1312**
- ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | **1314**
- Configuring the ANCP Agent | **1253**
- 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 | **1293**
- Preservation of CoS Shaping Across ANCP Agent Restarts | **1298**
- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | **1304**
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | **1306**
- Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | **1308**
- Verifying and Monitoring CoS for ANCP Subscribers | **1310**

Traffic Rate Reporting and Adjustment by the ANCP Agent

IN THIS SECTION

- Overview | **1293**
- Traffic Rate Adjustment | **1295**
- Recommended Traffic Shaping Rates | **1297**
- ANCP Agent Keepalives for CoS | **1298**

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 1295 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 1296 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	qos-adjust-adsl	–	adsl-bytes
ADSL2	qos-adjust-adsl2	–	adsl2-bytes
ADSL2+	qos-adjust-adsl2-plus	–	adsl2-plus-bytes
OTHER	qos-adjust-other	other-overhead-adjust	other-bytes
SDSL	qos-adjust-sdsl	sdsl-overhead-adjust	sdsl-bytes
VDSL	qos-adjust-vdsl	vdsl-overhead-adjust	vdsl-bytes

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 1297 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 ACI 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 | 1259

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 "dsl" on page 1818 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 sdsl-bytes bytes
user@host# set vdsl-bytes bytes
user@host# set vdsl2-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 sdsl-overhead-bytes bytes
user@host# set vdsl-overhead-bytes bytes
user@host# set vdsl2-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 sdsl-bonded-overhead-bytes bytes
user@host# set vdsl2-annex-q-bonded-overhead-bytes bytes
user@host# set vdsl2-annex-q-overhead-bytes bytes
user@host# set vdsl2-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 sdsl overhead-bytes bytes
user@host# set sdsl-bonded overhead-bytes bytes
user@host# set type tlv-value overhead-bytes bytes
user@host# set vdsl overhead-bytes bytes
user@host# set vdsl2 overhead-bytes bytes
user@host# set vdsl2-annex-q overhead-bytes bytes
user@host# set vdsl2-annex-q-bonded overhead-bytes bytes
user@host# set vdsl2-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 sdsl-overhead-adjust percentage
user@host# set vdsl-overhead-adjust percentage
user@host# set vdsl2-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 sdsl-overhead-adjust percentage
user@host# set vdsl-overhead-adjust percentage
user@host# set vdsl2-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 sdsl-bonded-overhead-adjust percentage
user@host# set vdsl2-annex-q-overhead-adjust percentage
user@host# set vdsl2-annex-q-bonded-overhead-adjust percentage
user@host# set vdsl2-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 sdsl overhead-adjust percentage
user@host# set sdsl-bonded overhead-adjust percentage
user@host# set type tlv-value overhead-adjust percentage
user@host# set vdsl overhead-adjust percentage
user@host# set vdsl2 overhead-adjust percentage
user@host# set vdsl2-annex-q overhead-adjust percentage
user@host# set vdsl2-annex-q-bonded overhead-adjust percentage
user@host# set vdsl2-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 | 1259

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# setqos-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 sdsl-bonded-total-adjust percentage
user@host# set vdsl2-annex-q-bonded-total-adjust percentage
user@host# set vdsl2-annex-q-total-adjust percentage
user@host# set vdsl2-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" on page 1818 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 adsl total-adjust percentage
user@host# set adsl2 total-adjust percentage
user@host# set adsl2-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 sdsl total-adjust percentage
user@host# set sdsl-bonded total-adjust percentage
user@host# set type tlv-value total-adjust percentage
user@host# set vdsl total-adjust percentage
user@host# set vdsl2 total-adjust percentage
user@host# set vdsl2-annex-q-bonded total-adjust percentage
user@host# set vdsl2-annex-q total-adjust percentage
user@host# set vdsl2-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 | 1310
- Action | 1310

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
```

Release History Table

Release	Description
19.3R1	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.
19.3R1	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.
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 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.
17.4R1	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.

RELATED DOCUMENTATION

ANCP Agent Neighbors and Operations | **1232**

Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | **1324**

Juniper Networks VSAs Supported by the AAA Service Framework | **19**

ANCP Agent and AAA

IN THIS SECTION

- ANCP Agent Interactions with AAA | **1312**
- ANCP TLVs Mapped to Juniper Networks and Broadband Forum Vendor-Specific Attributes | **1314**

Configuring AAA to Include Juniper Networks Access Line VSAs in RADIUS Messages | **1324**

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | **1325**

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 crates 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 1315 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	• 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	• 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 1320 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 1314 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 457 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 | **573**

RELATED DOCUMENTATION

ANCP Agent Neighbors and Operations | **1232**

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 | **1326**
- Verifying and Monitoring ANCP Neighbors | **1328**
- Clearing ANCP Neighbors | **1329**
- Verifying and Monitoring ANCP Subscribers | **1330**
- Clearing ANCP Subscribers | **1331**
- Clearing and Verifying ANCP Statistics | **1332**

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 | 1328
- Action | 1328

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 | 1329
- Action | 1329

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 | 1330
- Action | 1330

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 | 1331
- Action | 1331

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 | 1332
- Action | 1333

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 | 1232

Tracing ANCP Events for Troubleshooting

IN THIS SECTION

- Configuring the ANCP Trace Log Filename | 1334
- Configuring the Number and Size of ANCP Log Files | 1334
- Configuring Access to the ANCP Log File | 1335
- Configuring a Regular Expression for ANCP Messages to Be Logged | 1335
- Configuring the ANCP Tracing Flags | 1336
- Configuring the Severity Level to Filter Which ANCP Messages Are Logged | 1336

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 | 1232

8

PART

Diameter Base Protocol and its Applications

CHAPTER 12

Diameter Base Protocol and its Applications

IN THIS CHAPTER

- Diameter Base Protocol | **1338**
- Gx-Plus for Provisioning Subscribers | **1392**
- 3GPP Policy and Charging Control for Wireline Provisioning and Accounting | **1410**
- NASREQ for Authentication and Authorization | **1463**
- JSRC for Subscriber Provisioning and Accounting | **1467**
- JSRC and Subscribers on Static Interfaces | **1483**
- Monitoring and Management Diameter Information | **1499**
- Tracing Diameter Base Protocol Events for Troubleshooting | **1506**
- Troubleshooting Diameter Networks | **1510**
- Monitoring and Managing Static Subscriber Information | **1512**
- Tracing Static Subscriber Events for Troubleshooting | **1514**

Diameter Base Protocol

IN THIS SECTION

- Diameter Base Protocol Overview | **1339**
- Messages Used by Diameter Applications | **1342**
- Diameter AVPs and Diameter Applications | **1350**
- Configuring Diameter | **1373**
- Configuring the Origin Attributes of the Diameter Instance | **1374**
- Configuring Diameter Peers | **1374**
- Configuring the Diameter Transport | **1376**
- Configuring Diameter Network Elements | **1377**

- Example: Configure S6a Application | 1379

Diameter Base Protocol Overview

IN THIS SECTION

- Benefits of Using Diameter | 1341

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 1343 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	Request from Gx-Plus to the PCRF at subscriber login, logout, or update. 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. 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`. 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. An update request (CCR-U) message is sent when a usage threshold is reached. The CCR-U reports the actual usage for all

Table 73: Diameter Messages and Diameter Applications *(Continued)*

Diameter Message	Code	Application	Description
			statistics. The PCRF may return a CCA-U message that includes new monitoring thresholds, service activations, service deactivations. 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. 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. 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. 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.

Table 73: Diameter Messages and Diameter Applications *(Continued)*

Diameter Message	Code	Application	Description
Credit-Control-Answer (CCA)	272	Gx-Plus	Reply from the PCRF to a CCR message. 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. 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. A CCA-N is a response to a CCR-N.
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	Reply from router to a JSDR message; describes session information. The Result-Code AVP includes one of the following values, or an error value: • 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	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.
Juniper-Session-Event-Answer (JSEA)	8388628	Gx-Plus	Reply from PCRF to a JSER message.
Push-Profile-Request (PPR)	288	JSRC, PTSP	Request from the SAE to the router to activate or deactivate services for a subscriber.

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	Audit request from the PCRF to router to determine whether a specific subscriber is still present. The router updates the monitoring key and threshold values when they are received in the RAR.
Re-Auth-Answer (RAA)	258	Gx-Plus	Reply from router to a RAR message; indicates whether the subscriber is active. The Result-Code AVP includes one of the following values: • 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 1350 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	Specifies the number of seconds between each interim accounting update for this session. The router uses the following guidelines for interim accounting: • 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	Indicates whether a request completed successfully. Provides an error code if the request failed. The following classes are recognized by Diameter: • 1xxx—Informational • 2xxx—Success • 3xxx—Protocol errors • 4xxx—Transient errors • 5xxx—Permanent failures Unrecognized classes, which begin with numerals 6–9 or 0, are handled as permanent failures. JSRC and PTSP support the following values; all non-success values are treated as permanent failures: • 1001—DIAMETER MULTI ROUND AUTH • 2001—DIAMETER SUCCESS • 5002—DIAMETER UNKNOWN SESSION ID • 5012—DIAMETER UNABLE TO COMPLY JSRC also supports the following value, which is treated as a permanent failure:	Unsigned32

Table 74: Standard Diameter AVPs *(Continued)*

Attribute Number	Diameter AVP	Application	Description	Type
			• 3004—DIAMETER TOO BUSY; this is a transient condition, typically when the router already has a request in process for a specified subscriber. 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. • 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	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. If you change the product name, the router disconnects all existing connections to Diameter peers and reconnects using the new name.	UTF8String
277	Auth-Session-State	JSRC, NASREQ, PTSP	Indicates whether AAA session state is maintained. • 0—STATE MAINTAINED • 1—NO STATE MAINTAINED	Enumerated
279	Failed-AVP	NASREQ	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.	Grouped
281	Error-Message	NASREQ	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.	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 Diamter message is routed.	DiamIdent
295	Termination-Cause	JSRC, NASREQ, PTSP	Indicates the reason why a session was terminated on the access device. • 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: • 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	Contains the following sub-attributes that do no appear alone:
			• Subscription-Id-Type—(450) This subattribute has one of the following integer values:
			  • 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	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.	Grouped

Table 74: Standard Diameter AVPs *(Continued)*

Attribute Number	Diameter AVP	Application	Description	Type
480	Accounting-Record-Type	JSRC, PTSP	Specifies the type of account record for service accounting: • 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 1364 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: • 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: • 1—FULL-SYNC • 2—FAST-SYNC • 3—NO-STATE-TO-SYNC	Enumerated
2052	Juniper-Synchronization	2636	JSRC, PTSP	Describes the state of synchronization: • 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	Specifies whether to collect accounting data for the installed policy (service) when included in the Juniper-Policy-Install AVP: • 1—COLLECT_ACCT • 2—NOT_COLLECT_ACCT	Enumerated
2058	Juniper-State-ID	2636	JSRC, PTSP	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. **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.	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	Specifies the provisioning source for the session in CCR-N and JSDA messages: • 1—Local • 2—Remote	Enumerated
2102	Juniper-Provisioning-Descriptor	2636	Gx-Plus	Defines the group used in JSDA messages that includes the session ID, and optionally Juniper-Provisioning-Source and subscriber data.	Grouped
2103	Juniper-Event-Type	2636	Gx-Plus	Communicates the event type in JSER messages: • 1–Cold boot; all sessions are lost • 2—Warm boot; sessions are preserved • 3—Discovery requested by the operator • 4—*Are you there?* (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 JSDR and JSDA 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 JSDR and JSDA messages: • 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 JSDR and JSDA messages: • 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	If available in the subscriber's session database IPv6Prefix entry, this AVP is included in AAR provisioning request messages sent to the SAE. This AVP is used only when you enable JSRC dual-stack support.	IPv6Prefix
2201	Juniper-Framed-IPv6-Netmask	2636	JSRC	If available in the subscriber's session database IPv6Address entry, this AVP is included in AAR provisioning request messages sent to the SAE. This AVP is used only when you enable JSRC dual-stack support.	IPv6Address

Table 75: Juniper Networks Diameter AVPs *(Continued)*

Attribute Number	Diameter AVP	Vendor ID	Application	Description	Type
2202	Juniper-Agent-Circuit-Id	2636	JSRC	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. This AVP is used only when you enable JSRC dual-stack support.	OctetString
2203	Juniper-Agent-Remote-Id	2636	JSRC	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. This AVP is used only when you enable JSRC dual-stack support.	OctetString
2204	Juniper-Acct-IPv6-Input-Octets	2636	JSRC	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. This AVP is used only when you enable JSRC dual-stack support.	Unsigned64
2205	Juniper-Acct-IPv6-Output-Octets	2636	JSRC	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. This AVP is used only when you enable JSRC dual-stack support.	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	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. This AVP is used only when you enable JSRC dual-stack support.	Unsigned64
2207	Juniper-Acct-IPv6-Output-Pkts	2636	JSRC	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. This AVP is used only when you enable JSRC dual-stack support.	Unsigned64

Tekelec AVPs are used only for Gx-Plus. These AVPs have an enterprise number of 21274. Table 76 on page 1372 lists the Tekelec AVPs. These four variables are used to provide substitution values for user-defined CoS service variables.

Table 76: Tekelec Diameter AVPs

Attribute Number	Diameter AVP	Application	Description	Type
5555	Tekelec-Charging-Rule-Argument-Name	Gx-Plus	Defines the name of the service variable to be replaced.	OctetString
5556	Tekelec-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 1374
2. Configure the Diameter peers.

   See "Configuring Diameter Peers" on page 1374
3. (Optional) Configure the Diameter transport layer elements.

   See "Configuring the Diameter Transport" on page 1376
4. (Optional) Configure the Diameter network elements.

   See "Configuring Diameter Network Elements" on page 1377
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 1374.

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 | **1379**
- Overview | **1379**
- Configuration | **1380**
- Verification | **1389**

This example shows how to configure diameter-based authentication S6a application on your SRX series device to retrieve authentication information from the subscriber server.

Requirements

This example uses the following hardware:

- Any SRX Series device

Before you begin, read "Diameter Base Protocol Overview" on page 1339.

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 | 1380
- Configure Redundant Ethernet Interfaces | 1384
- Configure Security Zones and Security Policies to permit the S6a Diameter Application | 1387

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 | **1389**

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 shows the S6a application state and the statistics of the retrieved authentication information from the subscribed server.

Release History Table

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 | **1392**
- 3GPP Policy and Charging Control for Wireline Provisioning and Accounting | **1410**
- NASREQ for Authentication and Authorization | **1463**
- JSRC for Subscriber Provisioning and Accounting | **1467**

Gx-Plus for Provisioning Subscribers

IN THIS SECTION

- Gx-Plus for Provisioning Subscribers Overview | **1393**
- Understanding Gx-Plus Interactions Between the Router and the PCRF | **1395**
- Configuring Gx-Plus | **1404**
- Configuring the Gx-Plus Partition | **1405**
- Configuring Gx-Plus Global Attributes | **1406**
- Provisioning Subscribers with Gx-Plus | **1407**
- Disabling PCRF Control of a Subscriber Session | **1407**

Gx-Plus for Provisioning Subscribers Overview

IN THIS SECTION

- Benefits of Gx-Plus | **1395**

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 1393 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 | **1395**
- Fault Tolerance and Event Notification | **1398**
- PCRF-Generated Discovery | **1399**
- Subscriber Accounting | **1400**
- Subscriber Usage Thresholds | **1400**
- Subscriber Audit | **1404**
- Subscriber Logout | **1404**

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 40 to avoid overloading the PCRF. This limit reduces the possibility of losing requests. 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 1398 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.

Table 78: Router Events, Router Actions, and PCRF Actions *(Continued)*

Router Event	Router Action	PCRF Action
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 | 764

Configuring a Default Subscriber Service | 765

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 1405.
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 1406.
3. Configure Gx-Plus provisioning for subscribers.

   See "Provisioning Subscribers with Gx-Plus" on page 1407.
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 1407.
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 1373.

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

Release History Table

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 | **1338**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting | **1410**

NASREQ for Authentication and Authorization | **1463**

JSRC for Subscriber Provisioning and Accounting | **1467**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

IN THIS SECTION

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Gy File Backup Overview | **1438**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**
- Understanding Upstream and Downstream Messages for the PCRF | **1444**
- Configuring the OCS Partition | **1449**
- Configuring the PCRF Partition | **1455**
- Configuring OCS Global Parameters | **1462**

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

IN THIS SECTION

- Benefits of 3GPP Policy and Charging Control Architecture | **1412**

The 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) provides the unification of wireline provisioning and accounting for customers. Figure 50 on page 1411 shows the components of an overall 3GPP PCC architecture.

Figure 50: 3GPP PCC Architecture Overview

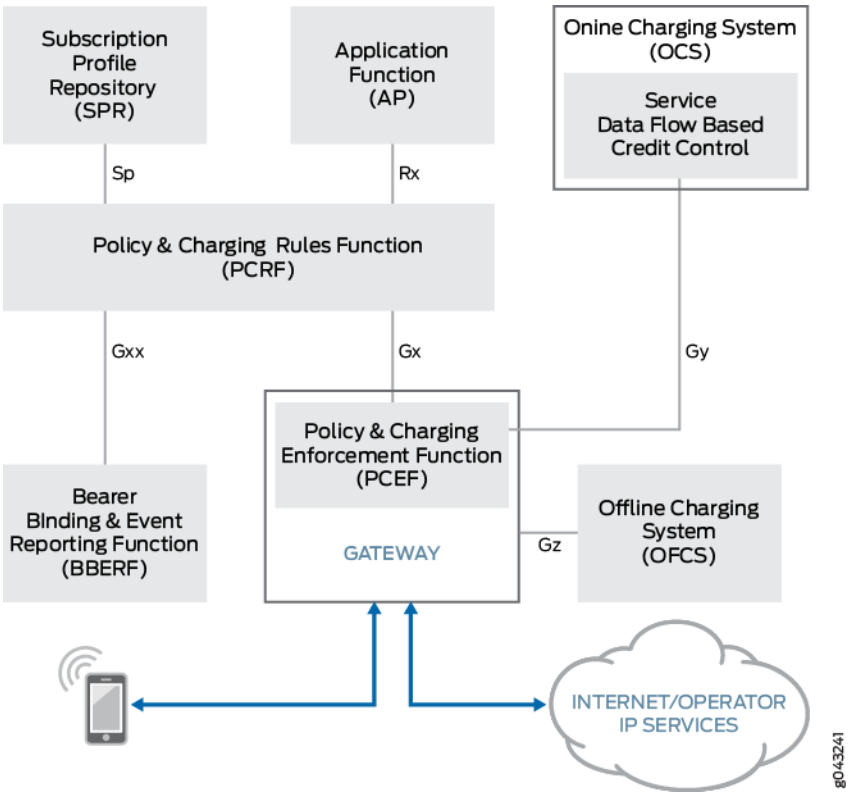

The four major components of the PCC architecture are:

- 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 1412 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 50 on page 1411 are:

- Application Function (AP)—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 | 1414
- Junos OS Environment | 1416

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 51 on page 1413.

Figure 51: 3GPP PCC Architecture Overview

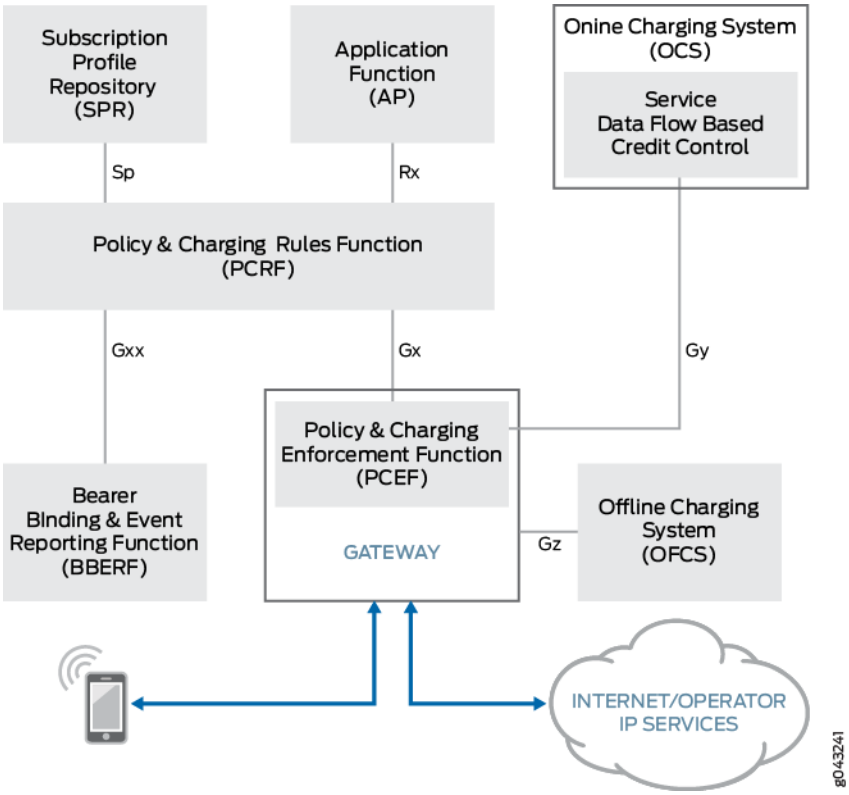

PCEF provides user traffic handling and quality of service (QoS) at the gateway, provides service data flow detection, and applies the rules received from the Policy Control and Charging Rules Function (PCRF). 3GPP defines Gx as the online policy protocol between the PCRF and the PCEF to provide control 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 1415 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 1416.

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 1417 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 | **1418**
- Subscriber Login Error Recovery | **1423**
- Subscriber Update | **1426**
- Subscriber Logout | **1427**
- Subscriber Disconnect | **1429**
- Connectivity Fault Recovery | **1430**

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 1420.

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 1421, 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 1421 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.

Table 84: CCA-GX-I Response Processing *(Continued)*

PCRF Local Decision Timeout	PCRF Local Decision	Result Category	Action
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.
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 1420).
- 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 1425 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.	The router does the following: • 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.	When the default provisioning completes, the router does the following: • 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.	The router does the following when no default services are configured: • 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. The router does the following when default services are configured: • Transitions to the `local-reinit-early` state. • Start provisioning the default services. When the default provisioning completes, the router does the following: • 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 1427.

Table 86: 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

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 1428.

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 1430.

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
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 | **1431**
- Intermediate Interrogation to the OCS | **1434**
- Final Interrogation to the OCS | **1436**
- Connectivity Fault Recovery | **1437**
- Abort Session Requests | **1437**

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 | **1438**
- Benefits of Gy File Backup | **1439**

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 stfp-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 | **1440**
- Update Interactions | **1441**
- Quota Expiration and Validity-Time Interactions | **1442**
- Connection and Monitoring Interactions | **1442**
- Logout Interactions | **1443**

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 52 on page 1440 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 52: 3GPP PCC Architecture Overview

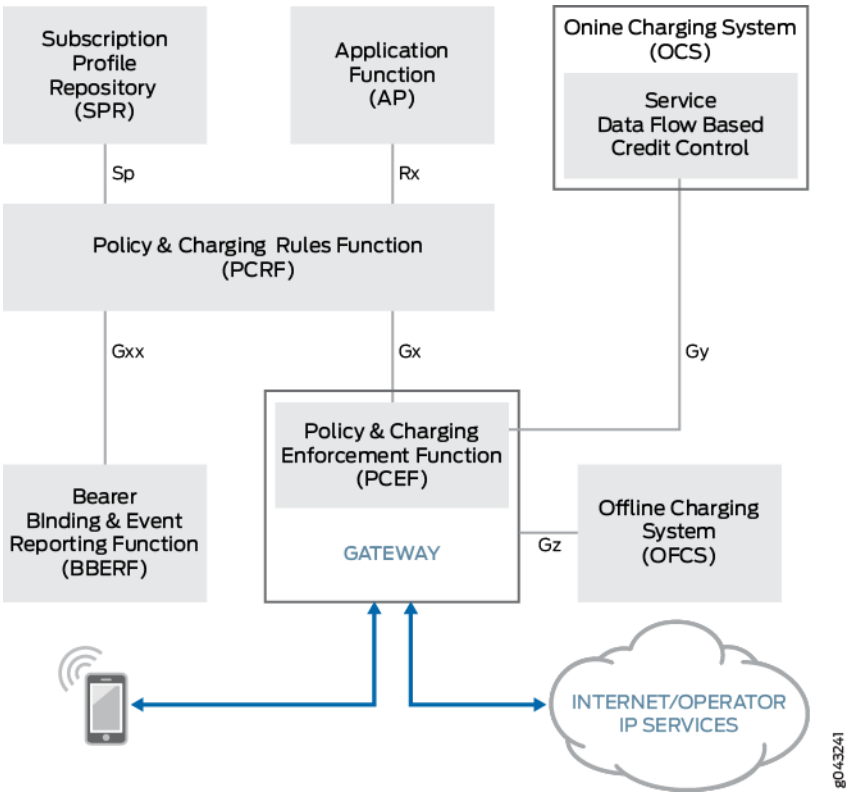

The PCRF can also push changes to rules applied to an existing session. However, modifications to rating groups is not supported. You are also required to set the `force-continue` statement at the `[edit access ocs partition` *`partition-name`*`] hierarchy level`.

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 | **1446**
- Common Downstream Messages | **1447**

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 1418 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 1446 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 1373.

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-mncggsn-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.

    a. (Optional) Include the underlying or physical interface name.

    ```
    [edit access ocs partition partition-name]
    user@host# set user-name-include base-interface-name
    ```

    b. (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
    ```

c. (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 stfp-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 1373.

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" on page 2163 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
```

Release History Table

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 | **1338**
- Gx-Plus for Provisioning Subscribers | **1392**
- NASREQ for Authentication and Authorization | **1463**
- JSRC for Subscriber Provisioning and Accounting | **1467**

NASREQ for Authentication and Authorization

IN THIS SECTION

- Diameter Network Access Server Application (NASREQ) | **1463**
- Configuring the Diameter Network Access Server Application (NASREQ) | **1465**

Diameter Network Access Server Application (NASREQ)

IN THIS SECTION

- Benefits of Using the Diameter NASREQ Protocol | **1465**

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 | 1338
- Gx-Plus for Provisioning Subscribers | 1392
- 3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1410
- JSRC for Subscriber Provisioning and Accounting | 1467

JSRC for Subscriber Provisioning and Accounting

IN THIS SECTION

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview | 1468
- Understanding JSRC-SAE Interactions | 1469
- JSRC Provisioning for Dual-Stack Subscribers | 1472
- JSRC Configuration Overview | 1476
- Configuring the JSRC Partition | 1477
- Assigning a Partition to JSRC | 1478
- Authorizing Subscribers with JSRC | 1478
- Provisioning Subscribers with JSRC | 1478
- Configuring JSRC for Dual-Stack Subscribers | 1479
- Excluding AVPs from Diameter Messages for JSRC | 1480
- Service Accounting with JSRC | 1480
- Configuring Service Accounting with JSRC | 1482

Juniper Networks Session and Resource Control (SRC) and JSRC Overview

IN THIS SECTION

- Benefits of Using JSRC | 1469
- Hardware Requirements for JSRC for Subscriber Access | 1469

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 | **1470**
- Subscriber Service Activation and Deactivation | **1470**
- Subscriber Resynchronization | **1471**
- Subscriber Session Terminated by the SAE | **1471**
- Statistics Collection and Reporting per Service Rule | **1472**
- Subscriber Logout | **1472**

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 | **1473**
- AA-Request Messages When Dual-Stack Support Configured | **1473**
- Accounting-Request Messages When Dual-Stack Support Configured | **1474**
- Network Family Activation and Deactivation Notification When Dual-Stack Support Configured | **1475**

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 1473 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.
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 1477.
2. Assign the JSRC partition.

   See "Assigning a Partition to JSRC" on page 1478.
3. Configure JSRC authorization for subscribers.

   See "Authorizing Subscribers with JSRC" on page 1478.
4. Configure JSRC provisioning for subscribers.

   See "Provisioning Subscribers with JSRC" on page 1478.
5. (Optional) Configure JSRC to exclude an AVP from Messages Sent to SAE.

   See "Excluding AVPs from Diameter Messages for JSRC" on page 1480.
6. Configure service accounting by JSRC.

   See "Configuring Service Accounting with JSRC" on page 1482.
7. Configure JSRC support for dual-stack subscribers.

   See "Configuring JSRC for Dual-Stack Subscribers" on page 1479.
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 1373.

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 1477

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
```

Release History Table

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 | 1338

Gx-Plus for Provisioning Subscribers | 1392

3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1410
NASREQ for Authentication and Authorization | 1463
JSRC and Subscribers on Static Interfaces | 1483

JSRC and Subscribers on Static Interfaces

IN THIS SECTION

- Subscribers on Static Interfaces Overview | 1483
- Subscribers over Static Interfaces Configuration Overview | 1487
- Example: Configuring Static Subscribers for Subscriber Access | 1488
- Specifying the Static Subscriber Global Access Profile | 1490
- Specifying the Static Subscriber Global Dynamic Profile | 1490
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers | 1491
- Configuring the Static Subscriber Global Authentication Password | 1492
- Configuring the Static Subscriber Global Username | 1492
- Creating a Static Subscriber Group | 1494
- Specifying the Static Subscriber Group Access Profile | 1495
- Specifying the Static Subscriber Group Dynamic Profile | 1495
- Specifying the Static Subscriber Group Service Profile | 1495
- Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group | 1496
- Configuring the Static Subscriber Group Authentication Password | 1497
- Configuring the Static Subscriber Group Username | 1497

Subscribers on Static Interfaces Overview

IN THIS SECTION

- Benefits of Subscribers on Static Interfaces | 1487

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 1490.
2. Specify the global dynamic profile that is instantiated when static subscribers log in.
   See "Specifying the Static Subscriber Global Dynamic Profile" on page 1490.
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 1491
4. Configure the global authentication password for static subscribers.
   See "Configuring the Static Subscriber Global Authentication Password" on page 1492.
5. Configure the global username for static subscribers.
   See "Configuring the Static Subscriber Global Username" on page 1492.
6. Configure a group of subscribers to share values different from the global configuration.
   See "Creating a Static Subscriber Group" on page 1494.
7. Specify the access profile for the static subscriber group.
   See "Specifying the Static Subscriber Group Access Profile" on page 1495.
8. Specify the dynamic profile for the static subscriber group.
   See "Specifying the Static Subscriber Group Dynamic Profile" on page 1495.
9. Specify the service profile for the static subscriber group.
   See "Specifying the Static Subscriber Group Service Profile" on page 1495.

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 1496.

11. Configure the authentication password for the static subscriber group.

    See "Configuring the Static Subscriber Group Authentication Password" on page 1497.

12. Configure the username for the static subscriber group.

    See "Configuring the Static Subscriber Group Username" on page 1497.

13. (Optional) Force a static subscriber to be logged out from an interface.

    See "Forcing a Static Subscriber to Be Logged Out" on page 1512.

14. (Optional) Enable an interface to accept static subscriber logins.

    See "Resetting the State of an Interface for Static Subscriber Login" on page 1512.

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 1513.

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 1513.

17. Configure trace options for troubleshooting the configuration.

    See "Tracing Static Subscriber Events for Troubleshooting" on page 1514.

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 jsrc;
        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" on page 1944 *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 `@`.

Release History Table

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 | 1467

Monitoring and Management Diameter Information

IN THIS SECTION

- Verifying Diameter Node, Instance, and Route Information | 1500
- Verifying and Managing Diameter Application Information | 1501
- Verifying and Managing Diameter Peer Information | 1503
- Verifying Diameter Network Element Information | 1505

Verifying Diameter Node, Instance, and Route Information

IN THIS SECTION

- Purpose | 1500
- Action | 1500

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 | 1501
- Action | 1502

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 | 1503
- Action | 1503

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 | 1505
- Action | 1505

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 | 1338

Tracing Diameter Base Protocol Events for Troubleshooting

IN THIS SECTION

- Configuring the Diameter Base Protocol Trace Log Filename | 1507
- Configuring the Number and Size of Diameter Base Protocol Log Files | 1507
- Configuring Access to the Diameter Base Protocol Log File | 1508
- Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged | 1508
- Configuring the Diameter Base Protocol Tracing Flags | 1509
- Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged | 1509

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 | 1338

Troubleshooting Diameter Networks

IN THIS SECTION

- Troubleshooting Diameter Network Configuration | 1510
- Troubleshooting Diameter Network Connectivity | 1511

Troubleshooting Diameter Network Configuration

IN THIS SECTION

- Problem | 1510
- Cause | 1511
- Solution | 1511

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 | 1511
- Cause | 1511
- Solution | 1512

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 | 1338

Monitoring and Managing Static Subscriber Information

IN THIS SECTION

- Forcing a Static Subscriber to Be Logged Out | 1512
- Resetting the State of an Interface for Static Subscriber Login | 1512
- Forcing a Group of Static Subscribers to Be Logged Out | 1513
- Resetting the State of an Interface Group for Static Subscriber Login | 1513
- Verifying Information about Subscriber Sessions on Static Interfaces | 1513

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 subscribes 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 | 1513
- Action | 1514

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 | 1483

Tracing Static Subscriber Events for Troubleshooting

IN THIS SECTION

- Configuring the Static Subscribers Trace Log Filename | 1515
- Configuring the Number and Size of Static Subscribers Log Files | 1515
- Configuring Access to the Static Subscribers Log File | 1516
- Configuring a Regular Expression for Static Subscriber Messages to Be Logged | 1516
- Configuring the Static Subscribers Tracing Flags | 1517
- Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged | 1517

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 | 1483

PART 9

Configuration Statements and Operational Commands

Configuration Statements | 1519

Operational Commands | 2537

CHAPTER 13

Configuration Statements

IN THIS CHAPTER

- aaa-logical-system (Domain Map) | **1534**
- aaa-routing-instance (Domain Map) | **1536**
- accept-max-tcp-connections (System Process) | **1537**
- accept-sdr (PCRF Partition) | **1539**
- access-identifier | **1540**
- access-line (Access-Line Rate Adjustment) | **1542**
- access-profile | **1557**
- access-profile (Extensible Subscriber Services Manager) | **1559**
- access-profile (Domain Map) | **1560**
- access-profile (Static Subscribers) | **1562**
- access-profile-name (Duplicate Accounting) | **1563**
- accounting (Access Profile) | **1565**
- accounting (Service Accounting) | **1566**
- accounting-backup-options (Access Profile) | **1568**
- accounting-order (Service Accounting) | **1569**
- accounting-stop-on-access-deny | **1571**
- accounting-stop-on-failure | **1572**
- active-leasequery (DHCP Relay Agent) | **1574**
- active-leasequery (DHCP Local Server) | **1577**
- active-server-group | **1579**
- actual-transit-statistics (Dynamic Profiles) | **1581**
- address (Diameter Peer) | **1582**
- address (Diameter Transport) | **1584**
- address-assignment (Address-Assignment Pools) | **1585**
- address-change-immediate-update | **1588**
- address-pool (Domain Map) | **1589**

- address-protection | **1591**
- address-ranges (Demux) | **1593**
- adjacency-timer | **1595**
- adsl-bytes | **1596**
- adsl2-bytes | **1598**
- adsl2-plus-bytes | **1600**
- advisory-options (Traffic Shaping) | **1602**
- aggregate-clients (DHCP Relay Agent) | **1603**
- aggregate-clients (Static Subscribers) | **1605**
- allow-active-leasequery (DHCP Local Server) | **1608**
- allow-bulk-leasequery (DHCP Local Server) | **1610**
- allow-leasequery (DHCP Local Server) | **1612**
- alternative-partition-name (OCS Partition) | **1614**
- always-write-giaddr | **1615**
- always-write-option-82 | **1617**
- ancp | **1618**
- ancp-speed-change-immediate-update (ANCP) | **1621**
- asymmetric-lease-time (DHCP Overrides) | **1622**
- asymmetric-prefix-lease-time (DHCP Overrides) | **1624**
- attempts (DHCP Local Server) | **1626**
- attributes (Access-Line Rate Adjustment) | **1628**
- attributes (RADIUS Attributes) | **1630**
- attributes (JSRC Attributes) | **1633**
- authentication (DHCP Local Server) | **1634**
- authentication (DHCP Relay Agent) | **1636**
- authentication (Static Subscribers) | **1638**
- authentication-order | **1640**
- authorization-order | **1642**
- authentication (Demux) | **1644**
- auto-configure (Demux) | **1646**
- auto-configure (IPv4) | **1647**
- auto-configure (IPv6) | **1650**

- autonomous (Dynamic Router Advertisement) | **1652**
- backup (OCS Partition) | **1653**
- bulk-leasequery (DHCP Relay Agent) | **1655**
- called-station-id (OCS Partition) | **1657**
- calling-station-id-format (Subscriber Management) | **1659**
- charging-id (OCS Partition) | **1661**
- charging-service-list | **1662**
- circuit-id (DHCP Relay Agent) | **1664**
- circuit-type (DHCP Local Server) | **1667**
- circuit-type (DHCP Relay Agent) | **1669**
- classification-key (DHCP Local Server) | **1671**
- classification-key (DHCP Relay Agent) | **1672**
- clear-on-abort (DHCP Local Server) | **1675**
- client-discover-match (DHCP Local Server) | **1677**
- client-discover-match (DHCP Relay Agent) | **1679**
- client-id (DHCP Local Server) | **1681**
- client-id (DHCP Relay Agent) | **1683**
- client-negotiation-match (DHCPv6 Local Server) | **1684**
- client-negotiation-match (DHCPv6 Relay Agent) | **1686**
- commit-interval | **1687**
- coa-immediate-update | **1689**
- coa-no-override service-class-attribute | **1690**
- concurrent-data-sessions | **1691**
- configuration-database (Enhanced Subscriber Management) | **1692**
- connect-actively | **1694**
- current-hop-limit (Dynamic Router Advertisement) | **1695**
- database-replication (Subscriber Session Database) | **1696**
- default-action (DHCP Relay Agent Option) | **1698**
- default-lifetime (Dynamic Router Advertisement) | **1699**
- delay-advertise (DHCPv6) | **1701**
- delay-authentication (DHCP Relay Agent) | **1704**
- delay-offer (DHCPv4) | **1705**

- delegated-pool (DHCP Local Server) | **1708**
- delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) | **1710**
- delimiter (DHCP Local Server) | **1711**
- delimiter (DHCP Relay Agent) | **1714**
- delimiter (Domain Map) | **1716**
- demux (Interfaces) | **1718**
- demux-options (All Demux Interfaces) | **1720**
- destination (Diameter Network Element) | **1721**
- destination-host | **1723**
- destination-host (Gx-Plus) | **1724**
- destination-host (OCS Partition) | **1725**
- destination-host (PCRF Partition) | **1726**
- destination-realm (JSRC) | **1728**
- destination-realm (Gx-Plus) | **1729**
- destination-realm (OCS Partition) | **1730**
- destination-realm (PCRF Partition) | **1732**
- dhcp-attributes (Address-Assignment Pools) | **1733**
- dhcp-local-server | **1740**
- dhcp-relay | **1752**
- dhcp-service | **1768**
- dhcpv6 (DHCP Local Server) | **1771**
- dhcpv6 (DHCP Relay Agent) | **1778**
- diameter | **1786**
- diameter-instance (JSRC) | **1788**
- diameter-instance (Diameter Applications) | **1789**
- dictionary | **1791**
- disable | **1792**
- disable (Extensible Subscriber Services Manager) | **1793**
- disable-relay | **1795**
- dne-origin (Diameter Network Element) | **1796**
- dns-server-address (Dynamic Profiles) | **1798**
- domain (Domain Map) | **1800**

- domain-name (DHCP Local Server) | **1801**
- domain-name (DHCP Relay Agent) | **1804**
- domain-name (Static Subscribers) | **1806**
- domain-name-server (Routing Instances and Access Profiles) | **1807**
- domain-name-server-inet (Routing Instances and Access Profiles) | **1809**
- domain-name-server-inet6 (Routing Instances and Access Profiles) | **1811**
- downstream-rate (Traffic Shaping) | **1812**
- draining (Diameter Applications) | **1814**
- draining-response-timeout (Diameter Applications) | **1815**
- drop (DHCP Relay Agent Option) | **1817**
- dsl (Access-Line Rate Adjustment) | **1818**
- dual-stack (DHCP Local Server Overrides) | **1824**
- dual-stack (DHCP Relay Agent Overrides) | **1825**
- dual-stack-group (DHCP Local Server) | **1827**
- dual-stack-group (DHCP Relay Agent) | **1830**
- dual-stack-interface-client-limit (DHCP Local Server and Relay Agent) | **1833**
- dualstack-support (JSRC) | **1834**
- duplication (Access Profile) | **1836**
- duplication-filter (Access Profile) | **1837**
- duplication-vrf (Duplicate Accounting) | **1839**
- dynamic-profile (Demux) | **1840**
- dynamic-profile (DHCP Local Server) | **1842**
- dynamic-profile (DHCP Relay Agent) | **1844**
- dynamic-profile (Domain Map) | **1846**
- dynamic-profile (Static Subscribers) | **1847**
- dynamic-profiles | **1849**
- enable | **1863**
- enable (Enhanced Subscriber Management) | **1864**
- equals (DHCP Relay Agent) | **1865**
- exceed-action | **1869**
- exclude (JSRC Attributes) | **1871**
- exclude (RADIUS Attributes) | **1872**

- excluded-address (Address-Assignment Pools) | **1881**
- excluded-range (Address-Assignment Pools) | **1882**
- external-authority | **1884**
- failover (System Process) | **1885**
- family (Address-Assignment Pools) | **1886**
- family-state-change-immediate-update | **1888**
- final-response-timeout (OCS Partition) | **1890**
- force-continue (OCS Partition) | **1891**
- forward-only (DHCP Relay Agent Option) | **1893**
- forward-only (DHCP Relay Agent) | **1894**
- forward-only-replies (DHCP Relay Agent) | **1897**
- forwarding (Diameter Network Element) | **1898**
- function (Diameter Network Element) | **1899**
- function (Diameter Route) | **1901**
- ggsn-address (OCS Partition) | **1903**
- ggsn-mcc-mnc (OCS Partition) | **1904**
- global (Gx-Plus) | **1906**
- global (OCS) | **1907**
- global (PCRF) | **1908**
- group (DHCP Local Server) | **1910**
- group (DHCP Relay Agent) | **1915**
- group (Static Subscribers) | **1921**
- gsmp-syn-timeout (ANCP) | **1923**
- gsmp-syn-wait (ANCP) | **1924**
- gx-plus (Gx-Plus) | **1926**
- host (Address-Assignment Pools) | **1927**
- host-name (DHCP Relay Agent) | **1929**
- host-name (DHCPv6 Relay Agent) | **1930**
- ietf-mode | **1931**
- immediate-update | **1933**
- include-ipv6 (Gx-Plus) | **1934**
- include-irb-and-l2 | **1935**

- include-option-82 (DHCP Local Server) | **1938**
- inet (Interfaces) | **1940**
- inet6 (Interfaces) | **1942**
- interface (DHCP Local Server) | **1944**
- interface (DHCP Relay Agent) | **1947**
- interface (Dynamic Router Advertisement) | **1950**
- interface (Static Subscriber Group) | **1952**
- interface (Static Subscriber Username) | **1954**
- interface-client-limit (DHCP Local Server) | **1955**
- interface-client-limit (DHCP Relay Agent) | **1958**
- interface-delete (Subscriber Management or DHCP Client Management) | **1960**
- interface-description (DHCP Local Server) | **1961**
- interface-description (DHCP Relay Agent) | **1963**
- interface-description-format | **1965**
- interface-name (DHCP Local Server) | **1967**
- interface-name (DHCP Relay Agent) | **1968**
- interface-mib (Enhanced Subscriber Management) | **1970**
- interface-set (ANCP) | **1971**
- interface-traceoptions (DHCP) | **1973**
- interfaces (ANCP) | **1975**
- interfaces (Static and Dynamic Subscribers) | **1977**
- interim-rate (Access) | **1984**
- ip-address-first | **1985**
- ip-can-type (PCRF Partition) | **1986**
- jsrc (JSRC) | **1988**
- jsrc (Access Profile) | **1989**
- jsrc-partition | **1991**
- layer2-unicast-replies | **1992**
- keep-incoming-circuit-id (DHCP Relay Agent) | **1993**
- keep-incoming-interface-id (DHCP Relay Agent) | **1995**
- keep-incoming-remote-id (DHCP Relay Agent) | **1996**
- leasequery (DHCP Relay Agent) | **1998**

- lease-time-threshold (DHCP Local Server and DHCP Relay Agent) | **2000**
- lease-time-validation (DHCP Local Server and DHCP Relay Agent) | **2002**
- limit | **2003**
- linked-pool-aggregation (Address-Assignment Pools) | **2004**
- local (Flat-File Access Profile) | **2006**
- local-decision (PCRF Partition) | **2008**
- local-server-group (DHCP Relay Agent Option) | **2011**
- location (DHCP Relay Agent) | **2013**
- location (DHCPv6 Relay Agent) | **2014**
- logical-interface-unit-range | **2015**
- logical-system (Diameter Peer) | **2017**
- logical-system (Diameter Transport) | **2018**
- logical-system-name (Static Subscribers) | **2020**
- logical-system-name (DHCP Local Server) | **2021**
- logical-system-name (DHCP Relay Agent) | **2023**
- logout-response-timeout (PCRF Partition) | **2025**
- ltv-syslog-interval (System Process) | **2026**
- mac-address (DHCP Local Server) | **2028**
- mac-address (DHCP Relay Agent) | **2030**
- maintain-subscriber (Subscriber Management) | **2031**
- managed-configuration (Dynamic Router Advertisement) | **2033**
- map (Domain Map) | **2034**
- max-advertisement-interval (Dynamic Router Advertisement) | **2037**
- max-data-sessions-per-subscriber | **2039**
- max-db-size (Enhanced Subscriber Management) | **2040**
- max-failures | **2043**
- max-outstanding-requests (Diameter Applications) | **2044**
- max-pending-accounting-stops (Access Profile) | **2046**
- max-withhold-time (Access Profile) | **2047**
- maximum-discovery-table-entries | **2049**
- maximum-helper-restart-time | **2050**
- maximum-subscribers | **2052**

- metric (Diameter Route) | **2053**
- min-advertisement-interval (Dynamic Router Advertisement) | **2054**
- multi-address-embedded-option-response (DHCP Local Server) | **2056**
- nas-port-extended-format | **2057**
- nas-port-extended-format (Interfaces) | **2060**
- nas-port-id-format (Subscriber Management) | **2062**
- nas-port-options (RADIUS Options) | **2065**
- nas-port-type (Subscriber Management) | **2067**
- nas-port-type (RADIUS Options) | **2069**
- nasreq (Diameter Application) | **2072**
- neighbor (Define ANCP) | **2074**
- network | **2076**
- network-element (Diameter Base Protocol) | **2077**
- network-services | **2079**
- no-bind-on-request (DHCP Relay Agent) | **2081**
- no-unsolicited-ra (Enhanced Subscriber Management) | **2083**
- no-vlan-interface-name | **2084**
- not-present (DHCP Relay Agent) | **2087**
- ocs (Diameter Applications) | **2090**
- on-demand-ip-address | **2093**
- on-demand-address-allocation | **2095**
- on-link (Dynamic Router Advertisement) | **2096**
- option-order (DHCP Relay Agent) | **2098**
- option-15 (DHCP Relay Agent) | **2100**
- option-16 (DHCP Relay Agent) | **2103**
- option-60 (DHCP Local Server) | **2105**
- option-60 (DHCP Relay Agent) | **2107**
- option-77 (DHCP Relay Agent) | **2110**
- option-82 (DHCP Relay Agent) | **2112**
- option-82 (DHCP Local Server Authentication) | **2114**
- option-82 (DHCP Local Server Pool Matching) | **2116**
- option-82 (Address-Assignment Pools) | **2117**

- option-match | **2119**
- option-number (DHCP Relay Agent Option) | **2121**
- options (Access Profile) | **2122**
- order | **2132**
- origin (Diameter Base Protocol) | **2134**
- other-bytes | **2135**
- other-overhead-adjust | **2137**
- other-stateful-configuration (Dynamic Router Advertisement) | **2139**
- overhead-accounting (ANCP) | **2140**
- override-chap-password | **2142**
- override-password (Domain Map) | **2143**
- overrides (DHCP Local Server) | **2144**
- overrides (DHCP Relay Agent) | **2148**
- overrides (Enhanced Subscriber Management) | **2150**
- parse-direction (Domain Map) | **2154**
- parse-order (Domain Map) | **2155**
- partition | **2157**
- partition (Gx-Plus) | **2158**
- partition (NASREQ Diameter Application) | **2159**
- partition (OCS) | **2161**
- partition (PCRF) | **2163**
- partition (s6a) | **2166**
- password (Static Subscribers) | **2168**
- password (DHCP Local Server) | **2170**
- password (DHCP Relay Agent) | **2172**
- pcrf (Diameter Applications) | **2174**
- peer (Diameter Base Protocol) | **2176**
- peer (Diameter Network Element) | **2178**
- peer-ip-address-optional | **2179**
- peer-origin (Diameter Peer) | **2181**
- pon (Access-Line Rate Adjustment) | **2182**
- pool (Address-Assignment Pools) | **2187**

- pool (DHCP Local Server Overrides) | **2190**
- pool-match-order | **2192**
- port (Diameter Peer) | **2194**
- pre-ietf-mode | **2195**
- preauthentication-order (Access Profile) | **2196**
- preferred-lifetime (Dynamic Router Advertisement) | **2198**
- prefix (DHCP Relay Agent) | **2199**
- prefix (Address-Assignment Pools) | **2201**
- prefix (Dynamic Router Advertisement) | **2203**
- priority (Diameter Peer) | **2204**
- profile (Access) | **2205**
- process-inform | **2213**
- protocol-master | **2215**
- protocols (Dynamic Profiles) | **2218**
- provisioning-order (Diameter Applications) | **2221**
- proxy-mode | **2223**
- qos-adjust | **2225**
- qos-adjust-adsl | **2227**
- qos-adjust-adsl2 | **2228**
- qos-adjust-adsl2-plus | **2230**
- qos-adjust-other | **2232**
- qos-adjust-sdsl | **2234**
- qos-adjust-vdsl | **2236**
- qos-adjust-vdsl2 | **2238**
- radius (Access Profile) | **2239**
- radius-disconnect (DHCP Local Server) | **2244**
- radius-flow-tap | **2246**
- radius-options (Access) | **2250**
- radius-options (Interfaces) | **2251**
- radius-server | **2253**
- range (Address-Assignment Pools) | **2259**
- rapid-commit (DHCPv6 Local Server) | **2261**

- reachable-time (Dynamic Router Advertisement) | **2262**
- realm-delimiter (Domain Map) | **2264**
- realm-parse-direction (Domain Map) | **2265**
- reauthenticate (DHCP Local Server) | **2267**
- reconfigure (DHCP Local Server) | **2270**
- redundancy (M:N Subscriber Redundancy) | **2272**
- relay-agent-interface-id (DHCP Local Server) | **2276**
- relay-agent-interface-id (DHCPv6 Relay Agent) | **2278**
- relay-agent-interface-id (DHCPv6 Relay Agent Username) | **2280**
- relay-agent-remote-id (DHCP Local Server) | **2281**
- relay-agent-remote-id (DHCPv6 Relay Agent) | **2283**
- relay-agent-remote-id (DHCPv6 Relay Agent Username) | **2285**
- relay-agent-subscriber-id (DHCP Local Server) | **2287**
- relay-agent-subscriber-id (DHCPv6 Relay Agent) | **2288**
- relay-option (DHCP Relay Agent) | **2290**
- relay-option-vendor-specific (dhcpv6) | **2292**
- relay-option-82 | **2293**
- relay-server-group (DHCP Relay Agent Option) | **2296**
- relay-source | **2298**
- remote-id (DHCP Relay Agent) | **2300**
- remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) | **2303**
- replace-ip-source-with (DHCP Relay Agent) | **2305**
- report-interface-descriptions (Access) | **2307**
- report-local-rule (PCRF Partition) | **2308**
- report-resource-allocation (PCRF Partition) | **2310**
- report-successful-resource-allocation (PCRF Partition) | **2312**
- request-max-tcp-connections (System Process) | **2314**
- request-rate (Access) | **2315**
- requested-ip-network-match (DHCP Local Server) | **2317**
- retransmit-timer (Dynamic Router Advertisement) | **2318**
- revert-interval (Access) | **2320**
- route (Diameter Network Element) | **2321**

- router-advertisement (Dynamic Profiles) | **2323**
- routing-instance (Diameter Peer) | **2324**
- routing-instance (Diameter Transport) | **2326**
- routing-instance-name (DHCP Local Server) | **2327**
- routing-instance-name (DHCP Relay Agent) | **2329**
- routing-instance-name (Static Subscribers) | **2331**
- s6a | **2333**
- sdsl-bytes | **2335**
- sdsl-overhead-adjust | **2336**
- send-acct-status-on-config-change (Access Profile) | **2338**
- send-dyn-subscription-indicator (PCRF Partition) | **2340**
- send-network-family-indicator (PCRF Partition) | **2342**
- send-origin-state-id (Diameter Applications) | **2344**
- send-release-on-delete (DHCP Relay Agent) | **2345**
- server-duid-type (DHCP Local Server) | **2347**
- server-group | **2348**
- server-id-override | **2350**
- server-response-time (DHCP Relay Agent) | **2352**
- service (Service Accounting) | **2354**
- service-context-id (OCS) | **2355**
- service-profile (DHCP Local Server) | **2357**
- service-profile (DHCP Relay Agent) | **2359**
- service-profile (Static Subscribers) | **2361**
- services (System Services) | **2362**
- session-limit-per-username (Access Profile) | **2370**
- session-options | **2372**
- sftp-backup (OCS Partition) | **2376**
- shmlog (Shared Memory Log) | **2378**
- short-cycle-protection (DHCP Local Server and Relay Agent) | **2382**
- smg-service (Enhanced Subscriber Management) | **2384**
- source-interface-set-at-login | **2385**
- stacked-vlan-ranges (RADIUS Options) | **2387**

- starts-with (DHCP Relay Agent Option) | **2389**
- static-subscribers (Dynamic Service Provisioning) | **2393**
- statistics (Access Profile) | **2395**
- statistics (Service Accounting) | **2397**
- strict (DHCP Local Server) | **2398**
- strip-domain (Domain Map) | **2400**
- strip-username (Domain Map) | **2401**
- sub-domain | **2402**
- subscriber (Access Profile) | **2408**
- subscriber-packet-idle-timeout | **2410**
- subscriber-management (Subscriber Management) | **2412**
- subscriber-profile | **2414**
- subscription-id-data-include (PCRF Partition) | **2416**
- subscription-id-type (PCRF Partition) | **2418**
- target-logical-system (Domain Map) | **2420**
- target-routing-instance (Domain Map) | **2422**
- terminate-code | **2424**
- timeout (DHCP Local Server) | **2426**
- timeout-grace (Access) | **2428**
- token (DHCP Local Server) | **2430**
- trace (DHCP Local Server) | **2432**
- trace (DHCP Relay Agent) | **2433**
- traceoptions (ANCP) | **2435**
- traceoptions (DHCP) | **2438**
- traceoptions (Diameter Base Protocol) | **2441**
- traceoptions (Extensible Subscriber Services Manager) | **2444**
- traceoptions (Enhanced Subscriber Management) | **2445**
- traceoptions (General Authentication Service) | **2449**
- traceoptions (Static Subscribers) | **2451**
- transport (Diameter Base Protocol) | **2454**
- transport (Diameter Peer) | **2455**
- traps | **2457**

- trigger (DHCP Local Server) | **2459**
- trio-flow-offload | **2461**
- trust-option-82 | **2462**
- tunnel-profile (Domain Map) | **2464**
- underlying-interface (ANCP) | **2465**
- unique-nas-port (Access) | **2466**
- unit | **2468**
- unit (Dynamic Profiles Standard Interface) | **2479**
- update-interval | **2483**
- update-interval (Service Accounting) | **2485**
- update-response-timeout (PCRF Partition) | **2487**
- upstream-rate (Traffic Shaping) | **2489**
- use-interface-description | **2490**
- username-include (Demux) | **2493**
- username-include (DHCP Local Server) | **2495**
- username-include (DHCP Relay Agent) | **2497**
- user-name-include (OCS Partition) | **2500**
- username-include (Static Subscribers) | **2503**
- use-option-82 | **2505**
- use-primary (DHCP Relay Agent) | **2507**
- use-underlying-interface-mac | **2509**
- use-vlan-id | **2510**
- use-vlan-id (DHCP Relay Agent) | **2512**
- user-prefix (DHCP Local Server) | **2514**
- user-prefix (DHCP Relay Agent) | **2516**
- user-prefix (Static Subscribers) | **2518**
- valid-lifetime (Dynamic Router Advertisement) | **2520**
- vdsl-bytes | **2521**
- vdsl-overhead-adjust | **2523**
- vdsl2-bytes | **2525**
- vdsl2-overhead-adjust | **2526**
- vendor-specific (DHCP Relay Agent) | **2528**

- violation-action (DHCP Local Server and DHCP Relay Agent) | **2530**
- vlan-ranges (RADIUS Options) | **2531**
- vrf-name (Duplicate Accounting) | **2533**
- wait-for-acct-on-ack (Access Profile) | **2535**

aaa-logical-system (Domain Map)

IN THIS SECTION

- Syntax | **1534**
- Hierarchy Level | **1534**
- Description | **1535**
- Default | **1535**
- Options | **1535**
- Required Privilege Level | **1535**
- Release Information | **1535**

Syntax

```
aaa-logical-system logical-system-name {
     aaa-routing-instance routing-instance-name;
}
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Configure a non-default logical system in which the `authd` daemon sends AAA requests for the domain map.

> **NOTE**: Subscriber management is supported in the default logical system only. The `aaa-logical-system` statement is for future extensions of subscriber management and is not supported in current Junos OS releases.

Default

Default logical system for the subscriber.

Options

`logical-system-name`—Name of the logical system.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 an AAA Logical System/Routing Instance in a Domain Map | 664

aaa-routing-instance (Domain Map)

IN THIS SECTION

- Syntax | 1536
- Hierarchy Level | 1536
- Description | 1536
- Default | 1536
- Options | 1537
- Required Privilege Level | 1537
- Release Information | 1537

Syntax

```
aaa-routing-instance (routing-instance-name | default);
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit access domain map domain-map-name aaa-logical-system logical-system-name]
```

Description

Configure the routing instance in which the `authd` daemon sends AAA requests for the domain map.

NOTE: Subscriber management is supported in the default logical system only. The `aaa-logical-system` statement, which appears in the CLI, is not supported in current Junos OS releases.

Default

Routing instance used for the subscriber context.

Options

routing-instance-name—Name of the routing instance.

`default`—The default routing instance.

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.

`default` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Specifying an AAA Logical System/Routing Instance in a Domain Map | 664

Domain Mapping Overview | 655

accept-max-tcp-connections (System Process)

IN THIS SECTION

- Syntax | 1538
- Hierarchy Level | 1538
- Description | 1538
- Options | 1538
- Required Privilege Level | 1538
- Release Information | 1538

Syntax

```
accept-max-tcp-connections max-tcp-connections;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Specify the number of simultaneous TCP connections that DHCP can use for bulk leasequery operations. This chassis-wide to enables the jdhcpd process to avoid exhaustion of TCP resources across the chassis.

NOTE: Use the `max-connections` option of the "allow-bulk-leasequery" on page 1610 statement to set the number of TCP connections allowed in the local server's logical system/routing instance.

Options

max-tcp-connections	Number of connections. • **Range:** 1 through 10 • **Default:** 5

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery | 812

Configuring and Using DHCP Bulk Leasequery | 814

accept-sdr (PCRF Partition)

IN THIS SECTION

- Syntax | 1539
- Hierarchy Level | 1539
- Description | 1539
- Required Privilege Level | 1539
- Release Information | 1540

Syntax

```
accept-sdr;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the service discovery requests for the vendor. If `accept-sdr` is configured for PCRF partition, the application and the vendor ID is advertised in the diameter CER request.

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 17.3.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444
- 3GPP Policy and Charging Control for Wireline Provisioning and Accounting | 1410

access-identifier

IN THIS SECTION

- Syntax | 1540
- Hierarchy Level | 1541
- Description | 1541
- Options | 1541
- Required Privilege Level | 1541
- Release Information | 1541

Syntax

```
access-identifier identifier-string;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-name],
[edit protocols ancp interfaces interface-set]
```

Description

Associate an access-loop circuit identifier (ACI) with the VLAN or set of VLANs that carry traffic to the subscriber using that access loop; identify a particular subscriber. This statement requires that the name of the interface or interface set is statically configured or deterministic. This means that it can be used with dynamic or static interface sets, VLAN-tagged interface sets, or static VLAN/VLAN demux interfaces.

Options

identifier-string—Unique identifier string for the access loop circuit; also configured on the access node.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | **1253**

Associating an Access Node with Subscribers for ANCP Agent Operations | **1256**

access-line (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax for Releases Earlier than Junos OS Release 19.3R1 | **1542**
- Syntax for Junos OS Release 19.3R1 and Higher Releases. | **1543**
- Hierarchy Level | **1546**
- Description | **1546**
- Options | **1547**
- Required Privilege Level | **1556**
- Release Information | **1556**

Syntax for Releases Earlier than Junos OS Release 19.3R1

```
access-line {
    adsl-overhead-bytes bytes;
    adsl-total-adjust percentage;
    adsl2-overhead-bytes bytes;
    adsl2-total-adjust percentage;
    adsl2-plus-overhead-bytes bytes;
    adsl2-plus-total-adjust percentage;
    gfast-bonded-overhead-adjust percentage
    gfast-bonded-overhead-bytes bytes
    gfast-bonded-total-adjust percentage
    gfast-overhead-adjust percentage
    gfast-overhead-bytes bytes
    gfast-total-adjust percentage
    hierarchical-access-network-detection;
    other-overhead-adjust percentage;
    other-overhead-bytes bytes;
    other-total-adjust percentage;
    sdsl-bonded-overhead-bytes  bytes
    sdsl-bonded-overhead-adjust  percentage
    sdsl-bonded-total-adjust  percentage
    sdsl-overhead-adjust percentage;
    sdsl-overhead-bytes bytes;
```

```
    sdsl-total-adjust percentage;
    vdsl-overhead-adjust percentage;
    vdsl-overhead-bytes bytes;
    vdsl-total-adjust percentage;
    vdsl2-annex-q-bonded-overhead-adjust  percentage
    vdsl2-annex-q-bonded-overhead-bytes bytes
    vdsl2-annex-q-bonded-total-adjust  percentage
    vdsl2-annex-q-overhead-adjust  percentage
    vdsl2-annex-q-overhead-bytes  bytes
    vdsl2-annex-q-total-adjust  percentage
    vdsl2-bonded-overhead-adjust  percentage
    vdsl2-bonded-overhead-bytes  bytes
    vdsl2-bonded-total-adjust  percentage
    vdsl2-overhead-adjust percentage;
    vdsl2-overhead-bytes bytes;
    vdsl2-total-adjust percentage;
}
```

Syntax for Junos OS Release 19.3R1 and Higher Releases.

```
access-line {
    attributes {
        preference (dsl | pon);
    }
    dsl {
        adsl {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2 {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2-plus {
            overhead-bytes bytes;
            total-adjust percent;
        }
        gfast {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
```

```
        }
        gfast-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        type tlv-value {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2 {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2-annex-q {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2-annex-q-bonded {
            overhead-adjust percent;
```

```
            overhead-bytes bytes;
            total-adjust percent;
        }
    }
    hierarchical-access-network-detection;
    pon {
        gpon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        twdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        type tlv-value {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        wdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xg-pon1 {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xgs-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure values to adjust data rates by a percentage of the actual data rate, or adjust encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes. Depending on the value, it may be reported to AAA, CoS, or both.

The actual (unadjusted) downstream and upstream data rates, line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

Adjustments are applied to all subscribers using access lines of the specific subscriber access line type:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following messages and AVPs:

- AVP 24, Tx Connect Speed (ICCN message)
- AVP 38, Rx Connect Speed (ICCN message)
- AVP 97, Connect Speed Update (CSUN message)

NOTE: Starting in Junos OS Release 19.3R1, the pre-existing adjustment options were renamed and placed in the new `dsl` stanza. The old DSL options are deprecated, but they redirect to the new location. The `hierarchical-access-network-detection` option is unchanged.

BEST PRACTICE: We recommend that you update your scripts to use the "dsl" on page 1818 statement when you upgrade to Junos OS Release 19.3R1 or higher releases. The redirect function will be supported for only a limited time.

Options

adsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

NOTE: This option replaces the `adsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL access line. The adjusted rate is reported only to AAA.

NOTE: This option replaces the `qos-adjust-adsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

adsl2-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `adsl2-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl2-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2 access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-adsl2` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

adsl2-plus-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2+ access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `adsl2-plus-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl2-plus-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2+ access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-adsl2-plus` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent

- **Default:** 100 percent

gfast-bonded-overhead-adjust *percentage*

Adjustment factor in percent that is applied to the downstream and upstream bonded data overhead rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

gfast-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell bonded overhead for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure. Specify G.fast bonded value in bytes.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

gfast-bonded-total-adjust *percentage*

Adjustment factor in percent that is globally applied to the downstream and upstream bonded data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

gfast-overhead-adjust *percentage*

Adjustment factor in percent that is applied to the downstream and upstream data overhead rates for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

gfast-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

gfast-total-adjust *percentage*

Adjustment factor in percent that is globally applied to the downstream and upstream data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

hierarchical-access-network-detection

Enable parsing of ANCP subscriber access loop attributes (TLVs) for backhaul line identifiers, to recognize when the access node references a logical interface set rather than an individual subscriber. If the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x03) string begins with a # sign, then the remainder of the string represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. The string is used as the name of a CoS Level 2 interface set that groups subscribers.

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

These TLVs can be parsed in ANCP messages or PPPoE IA tags in PADR messages.

- **Default:** Disabled, in case some users include an initial # character for some other purpose.

other-overhead-adjust *percentage*

Adjust the actual downstream rate for an access line of DSL type OTHER by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `other-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

other-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an access line of DSL type OTHER to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `other-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

other-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an access line of DSL type OTHER. The adjusted rate is reported only to AAA.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `qos-adjust-other` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

sdsl-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for an SDSL bonded access line by multiplying the rate by the specified percentage.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

sdsl-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL bonded access line. This value is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

sdsl-overhead-adjust *percentage*

Adjust the actual downstream rate for an SDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `sdsl-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

sdsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `sdsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-sdsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-vdsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 annex q bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-annex-q-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q bonded access line. The adjusted rate is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 annex q access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-annex-q-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q access line. The adjusted rate is reported to AAA.

vdsl2-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-bonded-overhead-bytes*bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes

- **Default:** 0 bytes

vdsl2-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 bonded access line. The adjusted rate is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl2-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

Number of bytes added to or subtracted from the actual downstream frame overhead.

> **NOTE**: This option replaces the `vdsl2-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-vdsl2` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

The following options added in Junos OS 18.2R1:

gfast-bonded-overhead-adjust	vdsl2-annex-q-bonded-overhead-adjust
gfast-bonded-overhead-bytes	vdsl2-annex-q-bonded-overhead-bytes
gfast-bonded-total-adjust	vdsl2-annex-q-bonded-total-adjust
gfast-overhead-adjust	vdsl2-annex-q-overhead-adjust
gfast-overhead-bytes	vdsl2-annex-q-overhead-bytes
gfast-total-adjust	vdsl2-annex-q-total-adjust
sdsl-bonded-overhead-bytes	vdsl2-bonded-overhead-bytes

sdsl-bonded-overhead-adjust	vdsl2-bonded-total-adjust
sdsl-bonded-total-adjust	vdsl2-bonded-total-adjust

hierarchical-access-network-detection option added in Junos OS 18.4R1.

attributes, dsl, and pon statements added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | **1306**
- Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | **1308**
- Configuring the ANCP Agent | **1253**

access-profile

IN THIS SECTION

- Syntax | **1558**
- Hierarchy Level | **1558**
- Description | **1558**
- Options | **1558**
- Required Privilege Level | **1558**
- Release Information | **1559**

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit],
[edit forwarding-options dhcp-relay]
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name]
[edit forwarding-options dhcp-relay dhcpv6]
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
[edit logical-systems logical-system-name routing-instances routing-instance-name]
[edit interfaces interface-name auto-configure vlan-ranges],
[edit interfaces interface-name auto-configure stacked-vlan-ranges],
[edit routing-instances routing-instances-name]
[edit system services dhcp-local-server]
[edit system services dhcp-local-server group group-name]
[edit system services dhcp-local-server dhcpv6]
[edit system services dhcp-local-server dhcpv6 group group-name]
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name]
```

Description

After you have created the access profile that specifies authentication and accounting parameters, you must specify where the profile is used. Authentication and accounting will not run unless you specify the profile. You can attach access profiles globally at the `[edit]` hierarchy level, or you can apply them to DHCP clients or subscribers, VLANs, or to a routing instance.

Options

profile-name—Name of the access profile that you configured at the `[edit access profile name]` hierarchy level.

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 9.1.

RELATED DOCUMENTATION

- RADIUS Servers and Parameters for Subscriber Access | **476**
- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | **703**
- *Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*

access-profile (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | **1559**
- Hierarchy Level | **1559**
- Description | **1560**
- Options | **1560**
- Required Privilege Level | **1560**
- Release Information | **1560**

Syntax

```
access-profile access-profile-name;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Define the access profile name for time accounting. In most cases, the information about the access profile on the RADIUS server is the same as the information about the access profile of a control session. This configuration is mandatory.

Options

access-profile-name	Name of the access profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

access-profile (Domain Map)

IN THIS SECTION

- Syntax | **1561**
- Hierarchy Level | **1561**
- Description | **1561**
- Options | **1561**
- Required Privilege Level | **1561**
- Release Information | **1561**

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Access profile that defines the AAA services and options for subscribers associated with the domain map.

Options

profile-name—Name of access 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 an Access Profile in a Domain Map | 662

access-profile (Static Subscribers)

IN THIS SECTION

- Syntax | 1562
- Hierarchy Level | 1562
- Description | 1562
- Options | 1563
- Required Privilege Level | 1563
- Release Information | 1563

Syntax

```
access-profile profile-name;
```

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]
```

Description

Specify the access profile that triggers AAA services 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 this statement overrides the global configuration.

Options

profile-name—Name of the static subscriber access profile.

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 | 1487
- Specifying the Static Subscriber Global Access Profile | 1490
- Specifying the Static Subscriber Group Access Profile | 1495

access-profile-name (Duplicate Accounting)

IN THIS SECTION

- Syntax | 1564
- Hierarchy Level | 1564
- Description | 1564
- Options | 1564
- Required Privilege Level | 1564
- Release Information | 1564

Syntax

```
access-profile-name [profile-name];
```

Hierarchy Level

```
[edit access profile profile-name accounting duplication-vrf]
```

Description

Specify up to five access profiles, all in the same nondefault VRF (LS:RI combination), each of which lists one or more RADIUS accounting servers to which duplication accounting information is sent.

Options

profile-name Name of an access profile that lists RADIUS accounting servers for duplicate reporting.

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 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases.

RELATED DOCUMENTATION

Understanding RADIUS Accounting Duplicate Reporting | 578

Configuring Authentication and Accounting Parameters for Subscriber Access | 550

accounting (Access Profile)

IN THIS SECTION

- Syntax | 1565
- Hierarchy Level | 1566
- Description | 1566
- Required Privilege Level | 1566
- Release Information | 1566

Syntax

```
accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    address-change-immediate-update;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    duplication-vrf {
        access-profile-name profile-name;
        vrf-name vrf-name;
    }
    immediate-update;
    order [accounting-method];
    send-acct-status-on-config-change
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access | 570

RADIUS Authentication and Accounting Basic Configuration | 550

accounting (Service Accounting)

IN THIS SECTION

- Syntax | 1567
- Hierarchy Level | 1567
- Description | 1567
- Required Privilege Level | 1567

- Release Information | 1567

Syntax

```
accounting {
    statistics (time | volume-time);
    update-interval minutes;
}
```

Hierarchy Level

```
[edit access profile profile-name service]
```

Description

Define the service accounting configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 14.2R1.

RELATED DOCUMENTATION

Configuring Service Accounting | 586

Configuring Per-Subscriber Session Accounting | 573

Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 583

accounting-backup-options (Access Profile)

IN THIS SECTION

- Syntax | 1568
- Hierarchy Level | 1568
- Description | 1568
- Required Privilege Level | 1568
- Release Information | 1569

Syntax

```
accounting-backup-options {
    max-pending-accounting-stops number;
    max-withhold-time hold-time;
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure options for backing up RADIUS accounting stop requests when all RADIUS accounting servers in the profile are offline.

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 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting | 591

accounting-order (Service Accounting)

IN THIS SECTION

- Syntax | 1569
- Hierarchy Level | 1569
- Description | 1569
- Default | 1570
- Options | 1570
- Required Privilege Level | 1570
- Release Information | 1571

Syntax

```
accounting-order (activation-protocol | local | radius);
```

Hierarchy Level

```
[edit access profile profile-name service]
```

Description

Specify the method that is used for reporting subscriber service accounting.

When you specify the activation-protocol method, the service accounting statistics are reported differently depending on how the service is activated:

- When the service is activated by the RADIUS server, the statistics are reported to the RADIUS server.
- When the service is activated by JSRC, the statistics are reported to JSRC.
- When the service is activated by the CLI configuration or a command—for example, through DHCP—the statistics are reported as determined by the flat-file profile configured at the [edit access profile *profile-name* local] hierarchy level. In this case, statistics are not reported if no flat-file profile is configured.

Default

activation-protocol

Options

activation-protocol Send service accounting reports by means of the application that activates services, such as JSRC, RADIUS, or the CLI.

local Send service accounting information to a flat file stored locally on the router. This method requires you to also configure the name of a flat-file profile at the [edit access profile *profile-name* local] hierarchy level. This profile determines which statistics and which nonstatistical parameters are collected for the service.

> **NOTE**: When you configure the local option, both volume and time statistics are collected for the service accounting sessions. In this case, you must not configure the volume-time option at the [edit access profile *profile-name* service accounting statistics] hierarchy level, or an error is generated when you commit the configuration.

radius Send service accounting reports to the RADIUS server by means of the RADIUS protocol.

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 11.4.

`local` option added in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Service Accounting with JSRC | 1482

Service Accounting with JSRC | 1480

Configuring Service Accounting in Local Flat Files

accounting-stop-on-access-deny

IN THIS SECTION

- Syntax | 1571
- Hierarchy Level | 1571
- Description | 1572
- Required Privilege Level | 1572
- Release Information | 1572

Syntax

```
accounting-stop-on-access-deny;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure RADIUS accounting to send an Acct-Stop message when the AAA server refuses a client request for access.

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 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration | 550

accounting-stop-on-failure

IN THIS SECTION

- Syntax | 1572
- Hierarchy Level | 1573
- Description | 1573
- Required Privilege Level | 1573
- Release Information | 1573

Syntax

```
accounting-stop-on-failure;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure RADIUS accounting to send an Acct-Stop message when a subscriber session has been successfully authenticated and authorized, but then fails before an Acct-Start message is sent. By default, an Acct-Stop message is sent only if an Acct-Start message has been exchanged with the accounting server.

Consider a situation where RADIUS address pools are used to assign IP/IPv6 addresses. After a subscriber session is successfully authenticated, the RADIUS server authorizes the session by assigning an IP address from the RADIUS address pool and conveying that address in the Framed-IP-Address attribute. If a negotiation failure occurs at this point, the session is terminated before activating. The Acct-Start message is never sent because it is initiated by session activation. By default, an Acct-Stop message cannot be sent because the Acct-Start is never sent. However, if the `acct-stop-on-failure` statement is configured, the negotiation failure causes the Acct-Stop message to be sent, which explicitly notifies the RADIUS server that the session is disconnected and that it can free the allocated IP address back to the pool.

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 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration | 550

active-leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1574
- Hierarchy Level | 1574
- Description | 1574
- Options | 1575
- Required Privilege Level | 1576
- Release Information | 1576

Syntax

```
active-leasequery {
    idle-timeout seconds;
    peer-address address;
    timeout seconds;
    topology-discover;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for active leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure parameters that the DHCP relay agent uses when sending DHCP active lease query messages to obtain lease information from the DHCP servers in the logical system/routing instance.

NOTE: 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.

Options

idle-timeout *seconds*

(Optional) Specify the number of seconds that pass with no activity before the TCP connection is terminated and restarted.

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; the connection is idle. If this 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.

- **Range:** 10 through 3600
- **Default:** 60

peer-address *address*

(Optional) Specify the address of a peer relay agent to synchronize binding information. You can configure up to five peer addresses. If the subscriber configuration and interface names are identical on peer relay agents, then active leasequery provides 1:1, chassis-level redundancy between peers for binding information.

NOTE: You can also achieve 1:1 chassis redundancy with an M:N redundancy configuration where:

- You configure redundancy for all DHCP subscribers on the chassis
- You configure a single chassis to back up all of the subscriber redundancy groups

timeout *seconds*

(Optional) Specify the number of seconds that TCP read/write operations can be blocked before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 120

topology-discover (Optional) Activate the topology discovery mechanism to discover peer access links for a redundancy group and their connections to the local access link. You configure this option on peer relay agents hosted on BNGs that are configured for M:N subscriber redundancy.

> **NOTE**: For dual-stack subscribers, you must configure active leasequery with topology discovery for both DHCPv4 and DHCPv6.

1. Each peer relay agent sends topology discovery request messages over the TCP connection to each peer relay agent.
2. The receiving peer returns a topology discovery response message with information about its corresponding access interface.
3. The querying relay agent then updates its translation table to map the remote (peer) access interface with its local access interface.

Each relay agent peer initiates topology discovery for all of its access interfaces. Each peer then updates its own translation table according the information received in topology discovery responses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

`topology-discover` option added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring and Using DHCP Active Leasequery | **818**

M:N Subscriber Redundancy on BNG Overview | **1173**

dhcp-relay | **1752**

active-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 1577
- Hierarchy Level | 1577
- Description | 1577
- Options | 1578
- Required Privilege Level | 1578
- Release Information | 1578

Syntax

```
active-leasequery {
             topology-discovery;
             peer-address {
                address;
               }
             timeout timeout;
             idle-timeout idle-timeout;
       }
```

Hierarchy Level

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

```
[edit system services dhcp-local-server dhcpv6]
```

Description

Subscriber redundancy on DHCP server focuses on subscriber synchronization between the peer servers using active leasequery. If you configure actiive leasequery for the DHCP server, the peer servers establishes TCP connection and sends active leasequery.

Options

idle-timeout *idle-timeout*
Configure the idle time out value in seconds for the connection between the peer servers.

- Range: 10 thru 3600 seconds
- Default: 60 seconds

peer-address *address*
Configure IPv4 or IPv6 address of the peer DHCP servers. You can configure maximum five IP addresses.

timeout *timeout*
Configure the TCP read/write block timeout vlue to terminate and restart the TCP connection.

- Range: 10 thru 3600 seconds
- Default: 60 seconds

topology-discovery
Enable the topology discovery option.

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 21.4R1.

RELATED DOCUMENTATION

M:N Subscriber Service Redundancy on DHCP Server | 1219

active-server-group

IN THIS SECTION

- Syntax | 1579
- Hierarchy Level | 1579
- Description | 1580
- Options | 1580
- Required Privilege Level | 1580
- Release Information | 1580

Syntax

```
active-server-group server-group-name <allow-server-change>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relaygroup group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Apply a DHCP relay agent configuration to the named group of DHCP server addresses. The server group itself is configured with the `server-group` statement. You can apply an active server group globally or for specific groups of interfaces, configured with the `group` statement. An active server group applied to an interface group overrides a global configuration.

Options

allow-server-change (Optional) (DHCPv4 only) Enable the relay agent to accept and forward a DHCP request (renew or rebind) ACK message to the client from any DHCP local server in the active server group. Starting in Junos OS Release 18.4R1, this option also applies to DHCP information request (DHCPINFORM) ACK messages.

- **Default:** Forward ACK messages from only the original binding server.

server-group-name Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

`allow-server-change` option added in Junos OS Release 16.2R1.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups | 858

Configuring Group-Specific DHCP Relay Options | 857

actual-transit-statistics (Dynamic Profiles)

IN THIS SECTION

- Syntax | 1581
- Hierarchy Level | 1581
- Description | 1581
- Default | 1582
- Required Privilege Level | 1582
- Release Information | 1582

Syntax

```
actual-transit-statistics;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-unit]
```

Description

Enable `actual-transit-statistics` to enable the collection and reporting of accurate accounting statistics for subscribers. The accurate statistics provide packet-level and byte-level counts for all transmit and receive traffic forwarded across the specified interface. The counts do not include overhead bytes, filtered traffic, drops, discards, or control traffic. These accurate statistics match the accounting statistics used by RADIUS. You can display these statistics with the `show subscribers id session-id accounting-statistics` and `show subscribers interface interface accounting-statistics` commands.

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.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI | **576**
- show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)
- *Dynamic Profiles Overview*

address (Diameter Peer)

IN THIS SECTION

- Syntax | **1583**
- Hierarchy Level | **1583**
- Description | **1583**

- Options | 1583
- Required Privilege Level | 1583
- Release Information | 1583

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Configure the IP address for a Diameter remote peer.

Options

ip-address—IPv4 or IPv6 address of remote Diameter peer. The address family must match that for the local Diameter transport connection.

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 9.6.

Support for IPv6 addresses added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

address (Diameter Transport)

IN THIS SECTION

- Syntax | 1584
- Hierarchy Level | 1584
- Description | 1584
- Options | 1584
- Required Privilege Level | 1585
- Release Information | 1585

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Configure the source (local) IP address for the Diameter local transport connection.

Options

ip-address—IPv4 or IPv6 address of remote Diameter peer. The address family must match that for the remote Diameter peer.

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 11.2.

Support for IPv6 addresses added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring Diameter | 1373

address-assignment (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1585
- Hierarchy Level | 1586
- Description | 1586
- Options | 1587
- Required Privilege Level | 1587
- Release Information | 1587

Syntax

```
address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
```

```
    neighbor-discovery-router-advertisement ndra-pool-name;
    pool pool-name {
        active-drain;
        family family {
            dhcp-attributes {
                protocol-specific attributes;
            }
            excluded-address ip-address;
            excluded-range name low minimum-value high maximum-value;
            host hostname {
                hardware-address mac-address;
                ip-address ip-address;
            }
            network ip-prefix/<prefix-length>;
            prefix ipv6-prefix;
            range range-name {
                high upper-limit;
                low lower-limit;
                prefix-length prefix-length;
            }
        }
        hold-down;
        link pool-name;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure address-assignment pools that can be used by different client applications.

NOTE: Support for subordinate statements is platform-specific. See individual statement topics for support information.

Options

abated-utilization Generate SNMP traps for DHCP address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

- **Values:** *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

abated-utilization-v6 Generate SNMP traps for DHCPv6 address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

- **Values:** *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

high-utilization Generate an SNMP trap when the DHCP address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

- **Values:** *percentage*—Percentage used to generate a trap. Range: 2 through 99.

high-utilization-v6 Generate an SNMP trap when the DHCPv6 address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

- **Values:** *percentage*—Percentage used to generate a trap. Range: 2 through 99.

neighbor-discovery-router-advertisement Configure the name of the address-assignment pool used to assign the router advertisement prefix.

- **Values:** *ndra-pool-name*—Name of the address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

L2TP LNS Inline Service Interfaces

Configuring an Address-Assignment Pool Used for Router Advertisements

address-change-immediate-update

IN THIS SECTION

- Syntax | 1588
- Hierarchy Level | 1588
- Description | 1588
- Default | 1589
- Required Privilege Level | 1589
- Release Information | 1589

Syntax

```
address-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send an Interim-Accounting message to the RADIUS server immediately after on-demand IPv4 allocation and de-allocation.

Changes to this setting take effect for new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Default

This functionality is disabled by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes | 1123

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | 1115

address-pool (Domain Map)

IN THIS SECTION

- Syntax | 1590
- Hierarchy Level | 1590
- Description | 1590
- Options | 1590
- Required Privilege Level | 1590
- Release Information | 1590

Syntax

```
address-pool pool-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Specify the address pool used to assign addresses to subscribers associated with the domain map.

Options

pool-name—Name of address pool.

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 an Address Pool in a Domain Map | 663

address-protection

IN THIS SECTION

- Syntax | 1591
- Hierarchy Level | 1591
- Description | 1591
- Options | 1592
- Required Privilege Level | 1592
- Release Information | 1593

Syntax

```
address-protection {
    reassign-on-match;
}
```

Hierarchy Level

```
[edit access],
[edit logical-systems logical-system-name routing-instances routing-instance-name access]
```

Description

Prevent IPv4 addresses and IPv6 prefixes from being assigned to more than one subscriber session when you use AAA to supply IPv4 addresses.

For IPv4:

If enabled, the router checks the following attributes received from external servers:

- *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.

For IPv6:

If enabled, the router 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 matches a prefix in an address pool, the prefix is taken from the pool, provided 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 exactly match the pool's prefix length, the authentication request is denied. If configured, the Acct-Stop message includes the cause for termination.

Options

reassign-on-match

Enable reassignment of an address from an existing subscriber to a new subscriber requesting that address. The address in use must not be part of a locally configured pool and address protection must be enabled. The request from the new subscriber is still rejected, but the existing subscriber is sent a disconnect request to begin the logout process. This enables the new subscriber to renegotiate and be assigned that IP address.

If the requested address is in a locally configured pool, the existing subscriber is not disconnected.

NOTE: This option is not supported for IPv6.

- **Default:** Rejects the address request from the new subscriber; the existing subscriber remains intact with the IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

`reassign-on-match` option added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Configuring Duplicate IPv4 Address Protection for AAA | 1158

Configuring Duplicate IPv6 Prefix Protection for Router Advertisement | 1051

address-ranges (Demux)

IN THIS SECTION

- Syntax | 1593
- Hierarchy Level | 1594
- Description | 1594
- Required Privilege Level | 1594
- Release Information | 1594

Syntax

```
address-ranges {
    authentication {
        password password-string;
        username-include {
            auth-server-realm realm-string;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            source-address;
```

```
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        network ip-address {
            range name {
                low lower-limit;
                high upper-limit;
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure]
```

Description

Specify the address range for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

adjacency-timer

IN THIS SECTION

- Syntax | 1595
- Hierarchy Level | 1595
- Description | 1595
- Options | 1595
- Required Privilege Level | 1596
- Release Information | 1596

Syntax

```
adjacency-timer seconds;
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Specify a value for the interval that the ANCP agent proposes during negotiation to establish an adjacency, for all neighbors or a specific neighbor. 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.

Options

seconds—Number of seconds between adjacency messages.

- **Range:** 1 through 25 seconds
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

- Configuring the ANCP Agent | 1253
- Specifying the Interval Between ANCP Adjacency Messages | 1257
- Configuring ANCP Neighbors | 1255

adsl-bytes

IN THIS SECTION

- Syntax | 1596
- Hierarchy Level | 1597
- Description | 1597
- Options | 1597
- Required Privilege Level | 1597
- Release Information | 1597

Syntax

```
adsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Configuring the ANCP Agent | **1253**

adsl2-bytes

IN THIS SECTION

- Syntax | **1598**
- Hierarchy Level | **1598**
- Description | **1599**
- Options | **1599**
- Required Privilege Level | **1599**
- Release Information | **1599**

Syntax

```
adsl2-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL2 access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl2-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

adsl2-plus-bytes

IN THIS SECTION

- Syntax | 1600
- Hierarchy Level | 1600
- Description | 1600
- Options | 1601
- Required Privilege Level | 1601
- Release Information | 1601

Syntax

```
adsl2-plus-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL2+ access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl2-plus-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

advisory-options (Traffic Shaping)

IN THIS SECTION

- Syntax | 1602
- Hierarchy Level | 1602
- Description | 1602
- Required Privilege Level | 1603
- Release Information | 1603

Syntax

```
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name],
[edit interfaces demux0 unit logical-unit-number],
[edit interfaces interface-name unit logical-unit-number]
```

Description

Specify a recommended shaping rate to be applied to downstream or upstream traffic on an interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | 1304

Configuring the ANCP Agent | 1253

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

aggregate-clients (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1603
- Hierarchy Level | 1604
- Description | 1604
- Options | 1605
- Required Privilege Level | 1605
- Release Information | 1605

Syntax

```
aggregate-clients (merge | replace);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dhcpv6 group  group-name dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name dynamic-profile profile-name],
[edit forwarding-options dhcp-relay group  group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group  group-name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group  group-name dynamic-profile profile-name]
```

Description

Specify that the router merge (chain) client attributes such as firewall filters and CoS attributes or replace them when multiple client sessions exist on the same underlying VLAN. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

Not supported for IP demux subscriber interfaces.

Options

`merge`—Aggregate multiple client attributes for the same subscriber (logical interface)

`replace`—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Options `merge` and `replace` introduced in Junos OS Release 9.5.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | **1752**

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring Group-Specific DHCP Relay Options | **857**

aggregate-clients (Static Subscribers)

IN THIS SECTION

- Syntax | **1606**
- Hierarchy Level | **1606**
- Description | **1606**

- Default | **1607**
- Options | **1607**
- Required Privilege Level | **1607**
- Release Information | **1607**

Syntax

```
aggregate-clients (merge | replace);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name system services static-subscribers dynamic-profile
profile-name],
[edit logical-systems logical-system-name system services static-subscribers group group-name
dynamic-profile profile-name],
[edit routing-instances routing-instances-name system services static-subscribers dynamic-profile
profile-name],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name dynamic-profile profile-name],
[edit system services static-subscribers dynamic-profile profile-name],
[edit system services static-subscribers group group-name dynamic-profile profile-name]
```

Description

Specify for all static subscribers or for a group of static subscribers that the router merge (chain) subscriber (client) attributes such as firewall filters and CoS attributes or replace them when multiple subscriber sessions exist on the same underlying VLAN. The group version of this statement overrides the global version.

NOTE: This statement is not supported for IP demux subscriber interfaces.

NOTE: This statement is not supported for enhanced subscriber management.

Default

By default, multiple subscribers cannot be on the same logical interface.

Options

`merge`—Aggregate the attributes of multiple subscribers for the logical interface.

`replace`—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface.

Required Privilege Level

system—To view this statement in the configuration.

system-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 | **1487**

Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers | **1491**

allow-active-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 1608
- Hierarchy Level | 1608
- Description | 1608
- Options | 1609
- Required Privilege Level | 1609
- Release Information | 1609

Syntax

```
allow-active-leasequery {
    idle-timeout seconds;
    timeout seconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to listen for, process, and respond to active leasequery requests received on TCP connections on TCP port 67 for DHCPv4 and on TCP port 547 for DHCPv6.

NOTE: Because active leasequery is an extension of bulk leasequery, you must also configure bulk leasequery before you configure active leasequery.

Options

idle-timeout *seconds*
(Optional) Specify the number of seconds that pass with no activity before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 60

timeout *seconds*
(Optional) Specify the number of seconds that TCP read/write operations can be blocked before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 120

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Configuring and Using DHCP Active Leasequery | 818

allow-bulk-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 1610
- Hierarchy Level | 1610
- Description | 1610
- Options | 1611
- Required Privilege Level | 1612
- Release Information | 1612

Syntax

```
allow-bulk-leasequery {
    max-connections number-of-connections;
    max-empty-replies number-of-replies;
    restricted-requestor;
    timeout seconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to listen for, process, and respond to bulk leasequery requests received on TCP connections on TCP port 67 for DHCPv4 and on TCP port 547 for DHCPv6.

Options

max-connections *number-of-connections*

Specify the maximum number of concurrent TCP connections allowed in the logical system/routing instance. This setting helps you manage the resources that the `jdhcpd` daemon uses for bulk leasequery operations in the logical system/routing instance. The number of connections you specify for a logical system/routing instances must be less than the number of connections you specify for the global `accept-max-tcp-connections` statement.

- **Range:** 1 through 10
- **Default:** 3

max-empty-replies *number-of-replies*

Specify the maximum number of empty replies that the DHCP local server sends to a specific requestor. When the maximum number is reached, the DHCP server closes the connection.

An empty reply is a response sent from the DHCP local server that contains no bindings or has an option status code error. Empty replies are often the response to an unauthorized requestor that has sent an invalid or incorrect query that produces no binding information. By limiting the number of empty replies that the DHCP local server can send, you prevent the connection from being consumed by an unauthorized or malicious requester, and free up the DHCP local server to support legitimate requestors.

- **Range:** 1 through 100
- **Default:** 5

restricted-requestor

Specify that the DHCP local server responds to a bulk leasequery request by sending the binding information to restricted requestors only. This ensures that the requestor is the originator of the binding.

Restricted requestors are determined as follows:

- For DHCPv4 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 sent during binding creation.

timeout *seconds*

Specify the number of seconds that a connection on the TCP socket can be idle before the DHCP local server closes the connection. Closing inactive connections enables the DHCP local server to more efficiently apply resources in support of active TCP connections.

- **Range:** 1 through 1000
- **Default:** 120

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Bulk Leasequery | **814**

allow-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | **1613**
- Hierarchy Level | **1613**
- Description | **1613**
- Options | **1613**
- Required Privilege Level | **1614**
- Release Information | **1614**

Syntax

```
allow-leasequery {
    restricted-requestor;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to support individual leasequery operations by listening on the UDP socket and responding to leasequery requests received on UDP connections.

Options

restricted-requestor Specify that the DHCP local server responds to an individual leasequery request by sending the binding information to restricted requestors only. This ensures that the requestor is the originator of the binding.

Restricted requestors are determined as follows:

- For DHCPv4 individual leasequery requests, the giaddr of the requestor must match the giaddr of the client.
- For DHCPv6 individual leasequery requests, the requestor's client ID in the bulk leasequery message must match the relay ID sent during binding creation.

> **NOTE**: The `restricted-requestor` statement is not supported for the DHCPv6 `allow-leasequery` statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery | 812

Configuring and Using DHCP Bulk Leasequery | 814

alternative-partition-name (OCS Partition)

IN THIS SECTION

- Syntax | 1614
- Hierarchy Level | 1614
- Description | 1615
- Required Privilege Level | 1615
- Release Information | 1615

Syntax

```
alternative-partition-name;
```

Hierarchy Level

```
[edit access ocs partition alternative-partition-name alternative-partition-name]
```

Description

Configure alternative partition name for OCS/GY. When configured, the specific partition name is used to forward diameter data along alternative data path in the specific cases.

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 17.3.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

always-write-giaddr

IN THIS SECTION

- Syntax | **1616**
- Hierarchy Level | **1616**
- Description | **1616**
- Required Privilege Level | **1616**
- Release Information | **1616**

Syntax

```
always-write-giaddr;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Overwrite the gateway IP address (giaddr) of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

dhcp-relay | 1752

always-write-option-82

IN THIS SECTION

- Syntax | 1617
- Hierarchy Level | 1617
- Description | 1618
- Required Privilege Level | 1618
- Release Information | 1618

Syntax

```
always-write-option-82;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

Extended DHCP Relay Agent Overview | 695

ancp

IN THIS SECTION

Syntax | 1619

- Hierarchy Level | 1620
- Description | 1620
- Required Privilege Level | 1621
- Release Information | 1621

Syntax

```
ancp {
    adjacency-loss-hold-time seconds;
    adjacency-timer seconds;
    gsmp-syn-timeout seconds;
    gsmp-syn-wait;
    interfaces {
        interface-set interface-set-name {
            access-identifier identifier-string;
            underlying-interface underlying-interface-name;
        }
        interface-name {
            access-identifier identifier-string;
        }
    }
    maximum-discovery-table-entries entry-number;
    maximum-helper-restart-time;
    neighbor ip-address {
        adjacency-loss-hold-time seconds;
        adjacency-timer;
        auto-configure-trigger interface interface-name;
        ietf-mode;
        maximum-discovery-table-entries entry-number;
        pre-ietf-mode;
    }
    pre-ietf-mode;
    qos-adjust {
        adsl-bytes bytes;
        adsl2-bytes bytes;
        adsl2-plus-bytes bytes;
        other-bytes bytes;
        other-overhead-adjust percentage;
```

```
        sdsl-bytes bytes;
        sdsl-overhead-adjust percentage;
        vdsl-bytes bytes;
        vdsl-overhead-adjust percentage;
        vdsl2-bytes bytes;
        vdsl2-overhead-adjust percentage;
    }
    qos-adjust-adsl adjustment-factor;
    qos-adjust-adsl2 adjustment-factor;
    qos-adjust-adsl2-plus adjustment-factor;
    qos-adjust-other adjustment-factor;
    qos-adjust-sdsl adjustment-factor;
    qos-adjust-vdsl adjustment-factor;
    qos-adjust-vdsl2 adjustment-factor;
    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;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Junos OS ANCP agent features.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

ancp-speed-change-immediate-update (ANCP)

IN THIS SECTION

- Syntax | 1621
- Hierarchy Level | 1621
- Description | 1622
- Required Privilege Level | 1622
- Release Information | 1622

Syntax

```
ancp-speed-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure AAA to generate immediate interim accounting updates to the RADIUS server in response to ANCP agent notifications of rate changes on subscriber access lines.

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 13.3.

RELATED DOCUMENTATION

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications | 1325

Configuring Per-Subscriber Session Accounting | 573

Configuring the ANCP Agent | 1253

asymmetric-lease-time (DHCP Overrides)

IN THIS SECTION

- Syntax | 1623
- Hierarchy Level | 1623
- Description | 1623
- Options | 1623
- Required Privilege Level | 1624
- Release Information | 1624

Syntax

```
asymmetric-lease-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6...]
```

Description

Set an asymmetric lease time for an address that is sent to the DHCP client instead of the actual lease time granted by the DHCP local server. The asymmetric lease time is shorter than the granted lease. Asymmetric leases require the DHCP client to renew the lease more frequently, at shorter intervals. If the client does not renew the short lease before it expires, the address is released sooner than if the original lease had been sent to the client, conserving address resources. Although you can configure the asymmetric lease on either the DHCP relay agent or the DHCP local server, in most situations, configuring the asymmetric lease on the relay agent is more useful.

Options

seconds Length of the short lease duration for the DHCPv4 client.

- **Range:** 600 through 65,534

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

DHCP Asymmetric Leasing Overview | 784

Configuring DHCP Asymmetric Leasing | 786

asymmetric-prefix-lease-time (DHCP Overrides)

IN THIS SECTION

- Syntax | 1624
- Hierarchy Level | 1625
- Description | 1625
- Options | 1625
- Required Privilege Level | 1625
- Release Information | 1625

Syntax

```
asymmetric-prefix-lease-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6...]
```

Description

Set an asymmetric lease time for an address that is sent to the DHCPv6 client instead of the actual lease time granted by the DHCP local server. The asymmetric lease time is shorter than the granted lease. Asymmetric leases require the DHCP client to renew the lease more frequently, at shorter intervals. If the client does not renew the short lease before it expires, the address is released sooner than if the original lease had been sent to the client, conserving address resources. Although you can configure the asymmetric lease on either the DHCP relay agent or the DHCP local server, in most situations, configuring the asymmetric lease on the relay agent is more useful.

Options

seconds Length of the short lease duration for the DHCPv6 client.

- **Range:** 600 through 65,534

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

DHCP Asymmetric Leasing Overview | 784

Configuring DHCP Asymmetric Leasing | 786

attempts (DHCP Local Server)

IN THIS SECTION

- Syntax | 1626
- Hierarchy Level | 1626
- Description | 1627
- Options | 1627
- Required Privilege Level | 1627
- Release Information | 1627

Syntax

```
attempts attempt-count;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
```

```
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure how many attempts are made to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces before reconfiguration is considered to have failed. A group configuration takes precedence over a DHCP local server configuration.

Options

attempt-count—Maximum number of attempts.

- **Range:** 1 through 10
- **Default:** 8

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Configuring Dynamic Reconfiguration Attempts for DHCP Clients | 875

attributes (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 1628
- Hierarchy Level | 1628
- Description | 1628
- Options | 1629
- Required Privilege Level | 1629
- Release Information | 1629

Syntax

```
attributes {
    preference (dsl | pon);
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Selects the set of access line attributes (TLVs) that is saved and processed when the router receives both DSL TLVs and PON TLVs in ANCP port status messages or in PPPoE-IA tags. The default line type is PON. You select the `dsl` options to prefer the DSL TLVs. For example, you might do this when the PON TLVs are unreliable, perhaps because of some issue with the OLT. The `preference` option has no effect when the BNG receives attributes of only one access line type from the OLT.

NOTE: The OLT might redundantly report the PON access line attributes both in PON TLVs and by overloading DSL TLVs. In this case, the DSL-Type TLV (0x91) is set to OTHER and PON rates for the access line are presented in the DSL TLVs, Actual-Net-Data-Rate-Upstream (0x81) and Actual-Net-Data-Rate-Downstream (0x82).

Options

preference dsl Specifies that the router saves and processes DSL TLVs from the DSL-Line-Attributes TLV 0x04. The router discards the PON-Access-Line-Attributes TLV (0x12) and the PON TLVs that it contains.

preference pon Specifies that the router saves and processes PON TLVs from the PON-Line-Attributes TLV 0x12. The router discards the DSL-Access-Line-Attributes TLV (0x04) and the DSL TLVs that it contains.

- **Default:** pon

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | **1306**
- Configuring the ANCP Agent | **1253**

attributes (RADIUS Attributes)

IN THIS SECTION

- Syntax | 1630
- Hierarchy Level | 1631
- Description | 1631
- Options | 1631
- Required Privilege Level | 1632
- Release Information | 1632

Syntax

```
attributes {
    exclude {
        attribute-name packet-type;
        standard-attribute number {
            packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
        }
        vendor-id id-number {
            vendor-attribute vsa-number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
| accounting-stop ];
            }
        }
    }
    ignore {
        dynamic-iflset-name;
        framed-ip-netmask;
        idle-timeout;
        input-filter;
        logical-system-routing-instance;
        output-filter;
        session-timeout;
        standard-attribute number;
        vendor-id id-number {
```

```
            vendor-attribute vsa-number;
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Specify how the router or switch processes RADIUS attributes.

Options

"exclude" on page 1872

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the `ignore` statement.

The options for this statement are explained separately. Click the linked statement for details.

ignore

Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. Standard attributes and VSAs received in RADIUS messages take precedence over internally provisioned attribute values. Ignoring the attributes enables your internally provisioned values to be used instead. Contrast this behavior with that provided by the `exclude` statement.

Starting in Junos OS Release 18.1R1, you can specify RADIUS standard attributes with the attribute number. You can specify vendor-specific attributes (VSAs) with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be ignored. The configuration has no effect if you can 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.

- Values:
  - dynamic-iflset-name—Ignore Juniper Networks VSA 26-130, Qos-Set-Name.
  - framed-ip-netmask—Ignore RADIUS attribute 9, Framed-IP-Netmask.
  - idle-timeout—Ignore RADIUS attribute 28, Idle-Timeout.
  - input-filter—Ignore Juniper Networks VSA 26-10, Ingress-Policy-Name.
  - logical-system-routing-instance—Ignore Juniper Networks VSA 26-1.
  - output-filter—Ignore Juniper Networks VSA 26-11, Egress-Policy-Name.
  - session-timeout—Ignore RADIUS attribute 27, Session-Timeout.
  - standard-attribute *number*—RADIUS standard attribute number supported by your platform. You can enclose multiple values in square brackets to specify a list of attributes. If you configure an unsupported attribute, that configuration has no effect. Range: 1 through 255.
  - vendor-attribute *vsa-number*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. You can enclose multiple values in square brackets to specify a list of VSAs. If you configure an unsupported VSA, that configuration has no effect. Range: 1 through 255.
  - vendor-id *id-number*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect. Range: 1 through 16777215.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

Standard and Vendor-Specific RADIUS Attributes | 3

attributes (JSRC Attributes)

IN THIS SECTION

- Syntax | 1633
- Hierarchy Level | 1633
- Description | 1633
- Required Privilege Level | 1634
- Release Information | 1634

Syntax

```
attributes {
    exclude {
        user-name [ authorization-request | provisioning-request ];
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name jsrc]
```

Description

Specify how the router processes Diameter attributes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.2.

RELATED DOCUMENTATION

Excluding AVPs from Diameter Messages for JSRC | 1480

Understanding JSRC-SAE Interactions | 1469

authentication (DHCP Local Server)

IN THIS SECTION

- Syntax | 1634
- Hierarchy Level | 1635
- Description | 1635
- Required Privilege Level | 1636
- Release Information | 1636

Syntax

```
authentication {
    password password-string;
    username-include {
        circuit-type;
```

```
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description  (device-interface | logical-interface);
        interface-name ;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

authentication (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1636
- Hierarchy Level | 1637
- Description | 1637
- Required Privilege Level | 1637
- Release Information | 1638

Syntax

```
authentication {
    password password-string;
    username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description  (device-interface | logical-interface);
```

```
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit 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 ...]
```

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the `[edit...dhcpv6]` hierarchy levels to configure DHCPv6 support.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 1752

Specifying Authentication Support | 833

authentication (Static Subscribers)

IN THIS SECTION

- Syntax | 1638
- Hierarchy Level | 1639
- Description | 1639
- Required Privilege Level | 1639
- Release Information | 1639

Syntax

```
authentication {
    password password-string;
    username-include {
        delimiter delimiter-character;
        domain-name domain-name;
        interface;
        logical-system-name;
        routing-instance-name;
        user-prefix user-prefix-string;
```

```
        vlan-tags;
    }
}
```

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]
```

Description

Specify the authentication parameters that trigger the Access-Request message to AAA 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 this statement overrides the global configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487

Configuring the Static Subscriber Global Authentication Password | 1492

Configuring the Static Subscriber Group Authentication Password | 1497

authentication-order

IN THIS SECTION

- Syntax | 1640
- Hierarchy Level | 1640
- Description | 1640
- Options | 1641
- Required Privilege Level | 1642
- Release Information | 1642

Syntax

```
authentication-order [ authentication-methods ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which AAA tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, AAA tries the authentication methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both `authentication-order` and `authorization-order` are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Starting in Junos OS Release 18.2R1, the `password` option can also be used to specify that local authentication and local authorization is attempted for individual subscribers that are configured with the `subscriber` statement at the `[edit access profile` *`profile-name`*`]` hierarchy level.

Options

authentication-methods Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination:

- `nasreq`—Verify subscribers using the Diameter-based Network Access Server Requirements (NASREQ) protocol.
- `none`—No authentication is performed. Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

  NOTE: Subscriber access management does not support the `none` option; authentication fails when this option is specified.

- `password`—Verify the client using the information configured at the `[edit access profile` *`profile-name`* `client` *`client-name`*`]` hierarchy level.

  Subscriber access management does not support the `password` option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, this option is used to enable local authentication and optionally local authorization for individual subscribers. Local authentication is typically used when you do not have external authentication and authorization servers. The password itself must be configured with the `subscriber` statement in the same access profile. Local authentication is performed when a subscriber logs in with a matching username; it succeeds if the subscribers login password matches the password in the profile.

  If you have external authentication and authorization servers, you can use local authentication as a backup authentication method. In this case, configure `password` other than first in the list of methods.

- `radius`—Verify the client using RADIUS authentication services.
- `s6a`—Verify subscribers using the Diameter-based s6a protocol.

- **Default:** `password`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`none` option added in Junos OS Release 11.2.

`nasreq` option added in Junos OS Release 16.1.

`s6a` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configuring CHAP Authentication with RADIUS

RADIUS Authentication and Accounting Basic Configuration | 550

Example: Configure S6a Application

authorization-order

IN THIS SECTION

- Syntax | 1643
- Hierarchy Level | 1643
- Description | 1643
- Options | 1643
- Required Privilege Level | 1644
- Release Information | 1644

Syntax

```
authorization-order [ authorization-methods ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which AAA tries different methods to verify that a client is authorized access to the router or switch. For each login attempt, AAA tries the authorization methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both `authentication-order` and `authorization-order` are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Options

authorization-methods Ordered list of methods to use for authorization attempts. The list includes one or more of the following methods in any combination:

- `jsrc`—Use the JSRC application in an SRC environment to request authorization from the SAE when verifying that a subscriber can access the router or switch.
  - When you configure both this option and `authentication-order`, AAA ignores the authentication order setting for DHCP subscribers. For non-DHCP subscribers, AAA ignores the authorization order and applies only the authentication order.
  - When you configure only this option, AAA applies the authorization order to both DHCP and non-DHCP subscribers.
- `nasreq`—Use the NASREQ application to communicate with a NASREQ server for authorization of any subscriber type as an alternative to RADIUS authorization.
- `none`—No authorization is performed. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

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 9.6.

`nasreq` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Authorizing Subscribers with JSRC | 1478

authentication (Demux)

IN THIS SECTION

- Syntax | 1644
- Hierarchy Level | 1645
- Description | 1645
- Required Privilege Level | 1645
- Release Information | 1645

Syntax

```
authentication {
    password password-string;
    username-include {
        auth-server-realm realm-string;
        delimiter delimiter-character;
```

```
        domain-name domain-name;
        interface-name;
        source-address;
        user-prefix user-prefix-string;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges]
```

Description

Specify the authentication parameters for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

auto-configure (Demux)

IN THIS SECTION

- Syntax | 1646
- Hierarchy Level | 1647
- Description | 1647
- Required Privilege Level | 1647
- Release Information | 1647

Syntax

```
auto-configure {
    address-ranges {
        authentication {
            password password-string;
            username-include {
                auth-server-realm realm-string;
                delimiter delimiter-character;
                domain-name domain-name;
                interface-name;
                source-address;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            network ip-address {
                range name {
                    low lower-limit;
                    high upper-limit;
                }
            }
        }
    }

}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet]
[edit interfaces interface-name unit unit-number demux inet6]
```

Description

Enable the configuration of dynamic, auto-sensed subscriber interfaces for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

auto-configure (IPv4)

IN THIS SECTION

- Syntax | 1648
- Hierarchy Level | 1648
- Description | 1648
- Options | 1649
- Required Privilege Level | 1649

- Release Information | **1649**

Syntax

```
auto-configure {
                    address-ranges {
                        session-timeout seconds;
                        dynamic-profile dynamic-profile {
                            network network-address;
                        }
                        authentication {
                            username-include {
                                delimiter;
                                domain-name;
                                user-prefix;
                                auth-server-realm;
                                interface-name;
                                source-address;
                            }
                        }
                    }
                }
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces junos-interface-ifd-name unit junos-interface-unit family
inet
```

Description

Defines the parameters for creating dynamic VLAN (DVLAN) interface on receiving the first VLAN packet from the client.

Options

dynamic-profile *dynamic-profile*
Name of the dynamic profile.

session-timeout *seconds*
Duration of the active dynamic IP subscriber in seconds. If you do not configure the `session-timeout` value, the system does not delete the dynamic IP subscriber unless an active DHCP subscriber is available.

- Range: 600 thru 6000

network *network-address*
IPv4 network address.

username-include
As a part of subscriber authentication, you can define the username with the following strings:

- `delimiter`-Delimiter or separator character
- `domain-name`-Domain name
- `user-prefix`-User defined prefix
- `auth-server-realm`-Authentication server realm name
- `interface-name`-Interface name
- `source-address`-Source address

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | **1227**

auto-configure (IPv6) | **1650**

auto-configure (IPv6)

IN THIS SECTION

- Syntax | 1650
- Hierarchy Level | 1651
- Description | 1651
- Options | 1651
- Required Privilege Level | 1651
- Release Information | 1652

Syntax

```
auto-configure {
                    address-ranges {
                        session-timeout seconds;
                        dynamic-profile dynamic-profile {
                            prefix prefix;
                        }
                        authentication {
                            username-include {
                                delimiter;
                                domain-name;
                                user-prefix;
                                auth-server-realm;
                                interface-name;
                                source-address;
                            }
                        }
                    }
                }
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces junos-interface-ifd-name unit junos-interface-unit family
inet6
```

Description

Defines the parameters for creating dynamic VLAN (DVLAN) interface on receiving the first VLAN packet from the client.

Options

dynamic-profile *dynamic-profile*
Name of the dynamic profile.

session-timeout *seconds*
Duration of the active dynamic IP subscriber in seconds. If you do not configure the `session-timeout` value, the system does not delete the dynamic IP subscriber unless an active DHCP subscriber is available.

- Range: 600 thru 6000

prefix *prefix*
IPv6 address prefix.

username-include
As a part of subscriber authentication, you can define the username with the following strings:

- `delimiter`-Delimiter or separator character
- `domain-name`-Domain name
- `user-prefix`-User defined prefix
- `auth-server-realm`-Authentication server realm name
- `interface-name`-Interface name
- `source-address`-Source address

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery | 1227

auto-configure (IPv4) | 1647

autonomous (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 1652
- Hierarchy Level | 1652
- Description | 1653
- Default | 1653
- Required Privilege Level | 1653
- Release Information | 1653

Syntax

```
(autonomous | no-autonomous);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-name
prefix prefix]
```

Description

Specify whether prefixes in the router advertisement messages are used for stateless address autoconfiguration:

- `autonomous`—Use prefixes for address autoconfiguration.
- `no-autonomous`—Do not use prefixes for address autoconfiguration.

Default

`autonomous`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

backup (OCS Partition)

IN THIS SECTION

- Syntax | 1654
- Hierarchy Level | 1654
- Description | 1654
- Options | 1654

Required Privilege Level | 1654
Release Information | 1655

Syntax

```
backup {
    limit;
    timeout seconds;
    overflow (deny-login | drop-oldest);
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Provides file backup for OCS data when both primary and alternative paths to the OCS are down.

Options

limit — The limit on the total number of backup entries. The login fails if the backup entries exceeds the prescribed limit.

timeout *seconds* — The timeout for backup operation.

overflow (deny-login | drop-oldest) — Controls action on the number of backup entries over limit. The values deny-login denies andy logins until the number of backup entries drops under backup-limit), and the value discard-oldest discards the oldest backup entries.

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 18.1.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Gy File Backup Overview | 1438
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

bulk-leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1655
- Hierarchy Level | 1656
- Description | 1656
- Options | 1656
- Required Privilege Level | 1657
- Release Information | 1657

Syntax

```
bulk-leasequery {
    attempts number-of-attempts;
    timeout seconds;
    trigger automatic;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for bulk leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure 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.

NOTE: You must also configure support on the relevant DHCP local servers with the "allow-leasequery" on page 1612 statement.

Options

attempts *number-of-attempts*

Specify the number of times the DHCP relay agent attempts to send DHCP bulk leasequery messages to the configured DHCP servers in the logical system/routing instance. DHCP relay agent resends the query message if the configured `timeout` value is reached, and either a confirmed reply or a reply from all configured DHCP servers has not been received. DHCP relay agent sends the subsequent messages only to the DHCP servers that have not replied to previous queries.

- Range:
  - (DHCPv4 bulk leasequery) 1 through 10
  - (DHCPv6 bulk leasequery) 1 through 720
- Default:
  - (DHCPv4 bulk leasequery) 6
  - (DHCPv6 bulk leasequery) 360

timeout *seconds*

Specify the number of seconds that DHCP relay agent waits before resending a leasequery message to the DHCP servers when all servers have not responded to a previous message.

- **Range:** 1 through 10
- **Default:** 10

trigger automatic (DHCPv6 only) Specify that the DHCPv6 relay agent always sends a bulk leasequery message whenever the jdhcp daemon is started or restarted and there are no bound subscribers in the session database. The automatic trigger updates the DHCP relay agent's binding information for clients associated with the requesting DHCPv6 relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Bulk Leasequery | 814

called-station-id (OCS Partition)

IN THIS SECTION

- Syntax | 1658
- Hierarchy Level | 1658
- Description | 1658
- Options | 1658
- Required Privilege Level | 1658
- Release Information | 1658

Syntax

```
called-station-id station-name;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the Called-Station-Id AVP value in all CCR-GY packets.

Options

station-name Called station name defined for the Called-Station-Id AVP to include in all CCR-GY packets.

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 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition | **1449**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gy Interactions Between the Router and the OCS | **1431**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

calling-station-id-format (Subscriber Management)

IN THIS SECTION

- Syntax | 1659
- Hierarchy Level | 1659
- Description | 1659
- Default | 1660
- Options | 1660
- Required Privilege Level | 1660
- Release Information | 1660

Syntax

```
calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the information that the router includes in the Calling-Station-ID (RADIUS IETF attribute 31) that is passed to the RADIUS server during authentication and accounting. You can include one or more optional values in any combination.

Default

The router displays the Calling-Station-ID set by the client.

Options

`agent-circuit-id`—Include the agent circuit identifier (ACI) string, which uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. The 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 PADI and PADR control packets (for PPPoE traffic).

`agent-remote-id`—Include the agent remote identifier (ARI) string, which identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The ARI string is stored in either the DHCP option 82 field of DHCP messages (for DHCP traffic), or in the DSL Forum Agent-Remote-ID VSA [26-2] of PPPoE PADI and PADR control packets (for PPPoE traffic).

`interface-description`—Include the interface description value.

`interface-text-description`—Include the interface text description.

`mac-address`—Include the MAC address of the source device for the subscriber.

`nas-identifier`—Include the NAS-identifier (RADIUS IETF attribute 32), which specifies the name of the NAS that orginated the authentication or accounting request.

`stacked-vlan`—Include the stacked VLAN tag value.

`vlan`—Include the VLAN tag value.

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 13.1.

`mac-address`, `interface-text-description`, `stacked-vlan`, and `vlan` options added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring a Calling-Station-ID with Additional Options | 490

charging-id (OCS Partition)

IN THIS SECTION

- Syntax | 1661
- Hierarchy Level | 1661
- Description | 1661
- Options | 1661
- Required Privilege Level | 1661
- Release Information | 1662

Syntax

```
charging-id number;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the 3GPP-Charging-Id AVP value in all CCR-GY messages.

Options

number 3GPP-Charging-Id AVP value to include in all CCR-GY messages.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

charging-service-list

IN THIS SECTION

- Syntax | 1662
- Hierarchy Level | 1662
- Description | 1663
- Options | 1663
- Required Privilege Level | 1663
- Release Information | 1663

Syntax

```
charging-service-list ocs;
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the list of charging services with which to communicate for the broadband Policy and Charging Enforcement Function (BPCEF).

NOTE: Currently, if you configure this statement, you must also configure the `provisioning-order` statement to `pcrf`.

Options

ocs Use Online Charging Services (OCS) as the list of charging services with which to communicate.

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 16.2.

RELATED DOCUMENTATION

- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

circuit-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1664
- Hierarchy Level | 1664
- Description | 1665
- Required Privilege Level | 1666
- Release Information | 1667

Syntax

```
circuit-id {
    include-irb-and-l2;
    keep-incoming-circuit-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-
option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
relay-option-82]
```

Description

Specify the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix, textual description, or VLAN tag.

NOTE: For Layer 3 interfaces, when you configure relay-option-82 only, the Agent Circuit ID is the default. If no VLAN tags are configured, then the default is the logical interface device (IFL) name. For integrated routing and bridging (IRB) interfaces, the default is the Layer 2 IFL name and bridge domain name.

The interface to bridge domain relationship may be implicit (the interface is mapped to the bridge domain by the system based on 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.

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs), stacked VLANs (S-VLANs), or bridge domains is as follows:

```
(fe | ge)-fpc/pic/port.subunit
```

NOTE: For remote systems, the *subunit* is required and is used to differentiate an interface for remote systems.

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:

```
(fe | ge)-fpc/pic/port:vlan-id
```

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:

```
(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

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
```

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`include-irb-and-l2` , `no-vlan-interface-name`, and `use-vlan-id` options added in Junos OS Release 14.1.

Release History Table

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `circuit-id` CLI statement in a stateless DHCP relay configuration. You can configure stateless DHCP relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay`] hierarchy level.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

Using DHCP Relay Agent Option 82 Information | 751

circuit-type (DHCP Local Server)

IN THIS SECTION

- Syntax | 1667
- Hierarchy Level | 1668
- Description | 1669
- Required Privilege Level | 1669
- Release Information | 1669

Syntax

```
circuit-type;
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit system services dhcp-local-server group *group-name* authentication username-include]

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

- Specifying Authentication Support

circuit-type (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1669
- Hierarchy Level | 1670
- Description | 1670
- Required Privilege Level | 1670
- Release Information | 1670

Syntax

```
circuit-type;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[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 ...]
```

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

classification-key (DHCP Local Server)

IN THIS SECTION

- Syntax | 1671
- Hierarchy Level | 1671
- Description | 1671
- Options | 1672
- Required Privilege Level | 1672
- Release Information | 1672

Syntax

```
classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-group
dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name ]
```

Description

Different mechanisms to identify a single household.

Options

circuit-id	Circuit-id as key.
mac-address	MAC address of client.
remote-id	Remote-id as key.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview | **1004**

classification-key (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **1673**
- Hierarchy Level | **1673**
- Description | **1674**
- Options | **1674**
- Required Privilege Level | **1674**
- Release Information | **1674**

Syntax

```
classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit bridge-domains name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name ddual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
```

```
[edit logical-systems name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-
group-name ],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
dual-stack-group dual-stack-group-name ],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dual-stack-group
dual-stack-group-name ],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group ],
[edit routing-instances name forwarding-options dhcp-relay dual-stack-group dual-stack-group-
name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-
stack-group dual-stack-group-name ],
[edit routing-instances name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-
group-name ],
[edit system services dhcp dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-
group-name ],
[edit vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ]
```

Description

Different mechanisms to identify a single household.

Options

circuit-id	Circuit-id as key
mac-address	MAC address of client
remote-id	Remote-id as key

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview | 1004

clear-on-abort (DHCP Local Server)

IN THIS SECTION

- Syntax | 1675
- Hierarchy Level | 1675
- Description | 1676
- Default | 1676
- Required Privilege Level | 1676
- Release Information | 1676

Syntax

```
clear-on-abort;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
```

```
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Delete all DHCP clients or only the DHCP clients serviced by the specified group of interfaces when reconfiguration fails; that is, when the maximum number of retry attempts have been made without success. A group configuration takes precedence over a DHCP local server configuration.

Default

Restores the original client configuration when reconfiguration fails.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Configuring Deletion of the Client When Dynamic Reconfiguration Fails | 876

client-discover-match (DHCP Local Server)

IN THIS SECTION

- Syntax | 1677
- Hierarchy Level | 1677
- Description | 1677
- Default | 1678
- Options | 1678
- Required Privilege Level | 1678
- Release Information | 1678

Syntax

```
client-discover-match <option60-and-option82 | incoming-interface>;
```

Hierarchy Level

```
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ... overrides],
[edit logical-systems logical-system-name system services dhcp-local-server ...overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server ...overrides]
```

Description

Configure the match criteria DHCP local server uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.

Default

By default, DHCP uses the `option60-and-option82` option.

Options

`incoming-interface` (Optional) Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

> **NOTE**: The `overrides client-discover-match incoming-interface` configuration deletes and replaces the existing binding when a new device connects. This action differs from the `overrides interface-client-limit 1` statement, which retains the existing binding and rejects the newly connected client.

`option60-and-option82` (Optional) Use option 60 and option 82 information to identify subscribers.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`incoming-interface` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

- Understanding Differences Between Legacy DHCP and Extended DHCP | **691**
- Overriding the Default DHCP Local Server Configuration Settings | **707**
- DHCP Auto Logout Overview | **879**
- Allowing Only One DHCP Client Per Interface | **863**

client-discover-match (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1679
- Hierarchy Level | 1679
- Description | 1679
- Default | 1680
- Options | 1680
- Required Privilege Level | 1680
- Release Information | 1680

Syntax

```
client-discover-match <option60-and-option82 | incoming-interface>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group ... overrides]
```

Description

Configure the match criteria DHCP relay uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.

Default

By default, DHCP uses the `option60-and-option82` option.

Options

`incoming-interface` (Optional) Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

> **NOTE**: The `overrides client-discover-match incoming-interface` configuration deletes and replaces the existing binding when a new device connects. This action differs from the `overrides interface-client-limit 1` statement, which retains the existing binding rejects the newly connected client.

`option60-and-option82` (Optional) Use option 60 and option 82 information to identify subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`incoming-interface` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

- Extended DHCP Relay Agent Overview | 695
- Overriding the Default DHCP Relay Configuration Settings | 709
- DHCP Auto Logout Overview | 879
- Allowing Only One DHCP Client Per Interface | 863

client-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 1681
- Hierarchy Level | 1681
- Description | 1682
- Options | 1682
- Required Privilege Level | 1682
- Release Information | 1682

Syntax

```
client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcp authentication username-include],
[edit system services dhcp-local-server dhcp group group-name authentication username-include],
```

```
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server dual-stack-group <group-name> authentication username-
include]
```

Description

Specify that the client identifier (DHCP option 61) is concatenated with the username during the subscriber authentication or client authentication process.

Options

exclude-headers — By default, all headers that are part of the client identifier format in option 61 are included in the username string used for RADIUS authentication. Configure the `exclude-headers` option to exclude the use of headers in the username string.

use-automatic-ascii-hex-encoding — By default, all components of the client identifier are converted to ASCII hex to encode the username. Configure the `use-automatic-ascii-hex-encoding` option to use ASCII hex encoding only if there are non-ASCII characters in the client identifier.

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.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

client-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1683
- Hierarchy Level | 1683
- Description | 1683
- Options | 1684
- Required Privilege Level | 1684
- Release Information | 1684

Syntax

```
client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...]
```

Description

Specify that the client identifier (DHCP option 61) is concatenated with the username during the subscriber authentication or client authentication process.

Options

exclude-headers

By default, all headers that are part of the client identifier format in option 61 are included in the username string used for RADIUS authentication. Configure the `exclude-headers` option to exclude the use of headers in the username string.

use-automatic-ascii-hex-encoding

By default, all components of the client identifier are converted to ASCII hex to encode the username. Configure the `use-automatic-ascii-hex-encoding` option to use ASCII hex encoding only if there are non-ASCII characters in the client identifier.

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.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

client-negotiation-match (DHCPv6 Local Server)

IN THIS SECTION

- Syntax | 1685
- Hierarchy Level | 1685
- Description | 1685

- Options | 1685
- Required Privilege Level | 1685
- Release Information | 1686

Syntax

```
client-negotiation-match incoming-interface;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides]
```

Description

Configure the match criteria the DHCPv6 local server uses to uniquely identify IPv6 subscribers or clients.

Options

incoming-interface Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Automatically Logging Out DHCPv6 Clients | 884
- Understanding Differences Between Legacy DHCP and Extended DHCP | 691
- Allowing Only One DHCP Client Per Interface | 863

client-negotiation-match (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 1686
- Hierarchy Level | 1686
- Description | 1687
- Options | 1687
- Required Privilege Level | 1687
- Release Information | 1687

Syntax

```
client-negotiation-match incoming-interface;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides]
```

Description

Configure the match criteria the DHCPv6 relay agent uses to uniquely identify IPv6 subscribers or clients.

Options

`incoming-interface` Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Automatically Logging Out DHCPv6 Clients | 884
- Extended DHCP Relay Agent Overview | 695
- Allowing Only One DHCP Client Per Interface | 863

commit-interval

IN THIS SECTION

- Syntax | 1688
- Hierarchy Level | 1688
- Description | 1688
- Options | 1688

Required Privilege Level | 1688
Release Information | 1688

Syntax

```
commit-interval interval;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Specify the interval at which Extensible Subscriber Services Manager issues requests for committing op script configurations. Requests that are received and processed within the interval are committed in a batch at the end of the interval. Requests that are processed after the interval are committed at the end of the next commit interval. When no operation script is to be committed, no request for committing operation script configurations is issued.

Options

interval Length of the interval.

- **Range:** 10 through 3600 seconds
- **Default:** 20 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Op Script Overview

coa-immediate-update

IN THIS SECTION

- Syntax | 1689
- Hierarchy Level | 1689
- Description | 1689
- Required Privilege Level | 1689
- Release Information | 1690

Syntax

```
coa-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send an Acct-Update message to the RADIUS accounting server immediately following a CoA operation.

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 10.0.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

Configuring Per-Subscriber Session Accounting | 573

coa-no-override service-class-attribute

IN THIS SECTION

- Syntax | 1690
- Hierarchy Level | 1690
- Description | 1690
- Required Privilege Level | 1691
- Release Information | 1691

Syntax

```
coa-no-override service-class-attribute;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

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.

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 11.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

Configuring Per-Subscriber Session Accounting | 573

concurrent-data-sessions

IN THIS SECTION

- Syntax | 1691
- Hierarchy Level | 1692
- Description | 1692
- Options | 1692
- Required Privilege Level | 1692
- Release Information | 1692

Syntax

```
concurrent-data-sessions max-session-number;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name enable service-name]
```

Description

Specify the maximum number of sessions that are concurrently enabled for the named service. The system randomly selects the number of sessions and enables the named service, whereas other sessions are not allotted the named service. This facilitates to increase the limit on the number of resources a service can use.

Options

max-session-number—Maximum number of sessions concurrently enabled for the named service.

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 11.4.

configuration-database (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | **1693**
- Hierarchy Level | **1693**
- Description | **1693**
- Required Privilege Level | **1693**
- Release Information | **1693**

Syntax

```
configuration-database {
    max-db-size size;
}
```

Hierarchy Level

```
[edit system]
```

Description

Enhanced subscriber management leverages system shared memory to improve performance and scaling. Since this memory is used by both the Junos configuration and enhanced subscriber management, you need to set an upper limit on the memory available to the Junos configuration, which in turn determines the allocation available for enhanced subscriber management. Starting in Junos OS Release 20.1R1, the upper limit applies to the Junos configuration database together with the schema database.

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 15.1R3.

RELATED DOCUMENTATION

> *Configuring Junos OS Enhanced Subscriber Management*
>
> *Junos OS Enhanced Subscriber Management Overview*

connect-actively

IN THIS SECTION

- Syntax | 1694
- Hierarchy Level | 1694
- Description | 1694
- Default | 1694
- Required Privilege Level | 1695
- Release Information | 1695

Syntax

```
connect-actively  {
    port port-number;
    transport transport-name;
}
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Define the destination port and transport connection used to establish active connections to Diameter peer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

Port 3868 and an automatically assigned local address are used to establish active connections to a peer.

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

current-hop-limit (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 1695
- Hierarchy Level | 1696
- Description | 1696
- Options | 1696
- Required Privilege Level | 1696
- Release Information | 1696

Syntax

```
current-hop-limit number;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Default value placed in the hop count field of the IP header for outgoing packets.

Options

number—Hop limit. A value of 0 means the limit is unspecified by this router.

- **Range:** 0 through 255
- **Default:** 64

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

database-replication (Subscriber Session Database)

IN THIS SECTION

Syntax | 1697

- Hierarchy Level | 1697
- Description | 1697
- Required Privilege Level | 1697
- Release Information | 1697

Syntax

```
database-replication {
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Define operations for subscriber management session database replication processes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Tracing Subscriber Management Session Database Replication Events for Troubleshooting

default-action (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 1698
- Hierarchy Level | 1698
- Description | 1699
- Required Privilege Level | 1699
- Release Information | 1699

Syntax

```
default-action {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay group group-name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option],
[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 ...]
```

Description

Specify the action DHCP relay agent takes when the option string in client traffic does not satisfy any match criteria or when no match criteria are configured.

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

The `local-server-group` option is not supported for DHCPv6 relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

default-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 1700
- Hierarchy Level | 1700
- Description | 1700
- Options | 1700

Required Privilege Level | 1700
Release Information | 1700

Syntax

```
default-lifetime seconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Lifetime associated with a default router.

Options

seconds—Default lifetime. A value of 0 means this router is not the default router.

- **Range:** Maximum advertisement interval value through 9000 seconds
- **Default:** Three times the maximum advertisement interval value

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

delay-advertise (DHCPv6)

IN THIS SECTION

- Syntax | 1701
- Hierarchy Level | 1702
- Description | 1702
- Options | 1703
- Required Privilege Level | 1703
- Release Information | 1703

Syntax

```
delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
        equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6 group
group-name interface interface-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6 group
group-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6
overrides]

[edit logical-systems name system services dhcp-local-server dhcpv6 group group-name interface
interface-name overrides]
[edit logical-systems name system services dhcp-local-server dhcpv6 group group-name overrides]
[edit logical-systems name system services dhcp-local-server dhcpv6 overrides]

[edit routing-instances name system services dhcp-local-server dhcpv6 group group-name interface
interface-name overrides]
[edit routing-instances name system services dhcp-local-server dhcpv6 group group-name overrides]
[edit routing-instances name system services dhcp-local-server dhcpv6 overrides]

[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name overrides]
[edit system services dhcp-local-server dhcpv6 group group-name overrides]
[edit system services dhcp-local-server dhcpv6 overrides]
```

Description

Delay sending a DHCPv6-advertise message to specified clients. The configured server sends the advertise message when the delay timer expires. If the client is already bound to another server when it receives this advertise message, the DHCPv6 server that delayed the message releases the client. You can change the delay value at any time, but the change applies only to clients from which a solicit is received after the change; the new value does not apply to clients for which the response is already delayed.

Configuring the delay on DHCPv6 local servers enables load-balancing among multiple local servers on the network. When a client sends a solicit message, the delay prevents more than one server from replying at the same time. The delay applies on a per-client basis. You can configure the clients affected based on DHCPv6 option 15 (user class identifier), option 16 (vendor class identifier), option 18 (interface identifier; ACI). or option 37 (ARI).

If the configured server receives a second solicit message from the client, that means that no other server has responded. In this case, the configured server immediately replies to the client. This behavior enables the server to act as a back-up for other servers on the network.

Options

ascii-string Value of the option expressed as an ASCII string.

based-on Specify the DHCPv6 option received in the solicit message that is compared with the configured value on the local server.

- `option-15`—DHCPv6 option 15, user class identifier.
- `option-16`—DHCPv6 option 16, vendor class identifier.
- `option-18`—DHCPv6 option 18, interface identifier, equivalent to the Agent Circuit ID.
- `option-37`—DHCPv6 option 37, the Agent Remote ID.

seconds Time delay between receiving DHCPv6 solicit message and responding to the client with an advertise message.

- **Range:** 1 through 30
- **Default:** 3

equals Specify that the received option value from the client must match the configured string.

hexadecimal-string Value of the option expressed as a hexadecimal.

not-equals Specify that the received option value from the client must not match the configured string.

starts-with Specify that the received option value from the client must start with the configured string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | 719

delay-authentication (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1704
- Hierarchy Level | 1704
- Description | 1704
- Required Privilege Level | 1705
- Release Information | 1705

Syntax

```
delay-authentication;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay ... overrides],
[edit forwarding-options dhcp-relay dhcpv6 ... overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... overrides]
```

Description

Delay authentication of subscribers until the DHCP client sends a Request packet. This conserves managed resources by delaying the authorization process and the creation of an entry in the subscriber database until the DHCP request processing phase.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Overriding the Default DHCP Relay Configuration Settings | 709

delay-offer (DHCPv4)

IN THIS SECTION

- Syntax | 1705
- Hierarchy Level | 1706
- Description | 1706
- Options | 1707
- Required Privilege Level | 1708
- Release Information | 1708

Syntax

```
delay-offer {
    based-on (option-60 | option-77 | option-82) {
        equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
```

```
        not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server group group-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server overrides]

[edit logical-systems name system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems name system services dhcp-local-server group group-name overrides]
[edit logical-systems name system services dhcp-local-server overrides]

[edit routing-instances name system services dhcp-local-server group group-name interface interface-name overrides]
[edit routing-instances name system services dhcp-local-server group group-name overrides]
[edit routing-instances name system services dhcp-local-server overrides]

[edit system services dhcp-local-server group group-name interface interface-name overrides]
[edit system services dhcp-local-server group group-name overrides]
[edit system services dhcp-local-server overrides]
```

Description

Delay sending a DHCP-offer message to specified clients. The configured server sends the offer message when the delay timer expires. If the client is already bound to another server when it receives this offer message, the DHCPv4 server that delayed the message releases the client. You can change the

delay value at any time, but the change applies only to clients from which an offer is received after the change; the new value does not apply to clients for which the response is already delayed.

Configuring the delay on DHCPv4 local servers enables load-balancing among multiple local servers on the network. When a client sends a discover message, the delay prevents more than one server from replying at the same time. The delay applies on a per-client basis. You can configure the clients affected based on DHCPv4 option 60 (vendor class identifier), option 77 (user class identifier), or option 82 (ACI and ARI).

If the configured server receives a second discover message from the client, that means that no other server has responded. In this case, the configured server immediately replies to the client. This behavior enables the server to act as a back-up for other servers on the network.

Options

ascii ***ascii-string*** Value of the option expressed as an ASCII string.

based-on Specify the DHCPv4 option received in the discover message that is compared with the configured value on the local server.

- `option-60`—DHCPv4 option 60, vendor class identifier.
- `option-77`—DHCPv4 option 77, the user class identifier.
- `option-82`—DHCPv4 option 82, the ACI/ARI.

delay-time ***seconds*** Time delay between receiving DHCPv4 discover message and responding to the client with an offer message.

- **Range:** 1 through 30
- **Default:** 3

equals Specify that the received option value from the client must match the configured string.

hexadecimal ***hexadecimal-string*** Value of the option expressed as a hexadecimal.

not-equals Specify that the received option value from the client must not match the configured string.

starts-with Specify that the received option value from the client must start with the configured string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers | 719

delegated-pool (DHCP Local Server)

IN THIS SECTION

- Syntax | 1708
- Hierarchy Level | 1709
- Description | 1709
- Options | 1709
- Required Privilege Level | 1709
- Release Information | 1709

Syntax

```
delegated-pool pool-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services system services dhcp-local-server
dhcpv6 ...],
[edit routing-instances routing-instance-name system services system services dhcp-local-server
dhcpv6 ...]
```

Description

Specify the address pool that assigns the IA_PD address. A pool specified by RADIUS VSA 26-161 takes precedence over the pool specified by this `delegated-pool` statement.

Options

pool-name	Name of the address-assignment pool.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation | 955

Overriding the Default DHCP Local Server Configuration Settings | 707

delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 1710
- Hierarchy Level | 1710
- Description | 1710
- Required Privilege Level | 1711
- Release Information | 1711

Syntax

```
delete-binding-on-renegotiation;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides]
```

Description

Configure DHCP to override the default behavior when the local server or relay agent receives a DHCPv4 Discover or DHCPv6 Solicit message while in a bound state. In this case, DHCP drops the message and it is not processed. 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.

A consequence of the override behavior is that the time to complete a DHCP negotiation is prolonged if the client begins negotiation before the existing client entry is expired. The delay can be up to several seconds.

The default behavior (this statement is not configured) is that DHCP maintains the client entry if it receives a Discover or Solicit message that has a client ID that matches the existing client. DHCP then processes the new message using the existing client entry and responds to the client with an Offer or Advertise message.

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.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCP Behavior When Renegotiating While in Bound State | 712

delimiter (DHCP Local Server)

IN THIS SECTION

Syntax | 1712

- Hierarchy Level | **1712**
- Description | **1713**
- Options | **1713**
- Required Privilege Level | **1713**
- Release Information | **1713**

Syntax

```
delimiter delimiter-character;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
```

```
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the character used as the delimiter between the concatenated components of the username.

Options

delimiter-character—Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.

- **Default:** . (period)

> **NOTE**: When you include the "interface-description " on page 1961 in the username, the delimiter must not be a character that is part of the interface description. For example, if the text description is configured as “Backbone connection/PHL01”, you cannot use the forward slash (/) as the delimiter.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

delimiter (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1714
- Hierarchy Level | 1714
- Description | 1715
- Options | 1715
- Required Privilege Level | 1715
- Release Information | 1716

Syntax

```
delimiter delimiter-character;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
```

```
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify the character used as the delimiter between the concatenated components of the username. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

delimiter-character—Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.

- **Default:** . (period)

> **NOTE**: When you include the "interface-description " on page 1961 in the username, the delimiter must not be a character that is part of the interface description. For example, if the text description is configured as “Backbone connection/PHL01”, you cannot use the forward slash (/) as the delimiter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

delimiter (Domain Map)

IN THIS SECTION

- Syntax | 1716
- Hierarchy Level | 1717
- Description | 1717
- Default | 1717
- Options | 1717
- Required Privilege Level | 1717
- Release Information | 1717

Syntax

```
delimiter [delimiter-character];
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the characters that the router uses to separate usernames from domain names.

Default

The `@` character.

Options

delimiter-character—One or more characters used as delimiters. You can specify a maximum of eight delimiters. You cannot use the semicolon (;) as a delimiter. Do not include spaces between characters.

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 Domain and Realm Name Delimiters | **668**

Configuring Domain and Realm Name Usage for Domain Maps | **667**

demux (Interfaces)

IN THIS SECTION

- Syntax | 1718
- Hierarchy Level | 1719
- Description | 1719
- Required Privilege Level | 1720
- Release Information | 1720

Syntax

```
demux {
    inet {
        address source;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }
```

```
        }
    }
    inet6 {
        address source;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }

        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number]
```

Description

Configure demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

demux-options (All Demux Interfaces)

IN THIS SECTION

- Syntax | 1720
- Hierarchy Level | 1721
- Description | 1721
- Required Privilege Level | 1721
- Release Information | 1721

Syntax

```
demux-options {
    use-underlying-interface-mac {
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure demultiplexing (demux) interface options for all demux interfaces.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | 1052

destination (Diameter Network Element)

IN THIS SECTION

- Syntax | 1722
- Hierarchy Level | 1722
- Description | 1722
- Options | 1722
- Required Privilege Level | 1722

Release Information | 1722

Syntax

```
destination realm realm-name <host hostname>;
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding route dne-route-name]
```

Description

Associate the route with all hosts of the specified realm or with a specific host of the specified realm. Together with the function and metric, defines a route reachable through a Diameter network element.

Options

`host` *hostname*—(Optional) Name of the destination host associated with the route.

`realm` *realm-name*—Name of the destination realm associated with the route.

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

destination-host

IN THIS SECTION

- Syntax | 1723
- Hierarchy Level | 1723
- Description | 1723
- Options | 1723
- Required Privilege Level | 1723
- Release Information | 1724

Syntax

```
destination-host hostname
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Configure the host on which the SAE application resides.

Options

hostname—Host on which the SAE is installed.

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 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Configuring the JSRC Partition | 1477

destination-host (Gx-Plus)

IN THIS SECTION

- Syntax | 1724
- Hierarchy Level | 1724
- Description | 1724
- Options | 1725
- Required Privilege Level | 1725
- Release Information | 1725

Syntax

```
destination-host hostname;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
```

Description

Configure the host on which the PCRF application resides.

Options

hostname—Host on which the PCRF is installed.

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 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus | **1404**

Configuring the Gx-Plus Partition | **1405**

destination-host (OCS Partition)

IN THIS SECTION

- Syntax | **1725**
- Hierarchy Level | **1726**
- Description | **1726**
- Options | **1726**
- Required Privilege Level | **1726**
- Release Information | **1726**

Syntax

```
destination-host ocs-hostname;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the Destination-Host AVP value used in the CCR-GY-I message.

Options

ocs-hostname—Value of the Destination-Host AVP to be used in the CCR-GY-I message.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

destination-host (PCRF Partition)

IN THIS SECTION

- Syntax | **1727**

- Hierarchy Level | 1727
- Description | 1727
- Options | 1727
- Required Privilege Level | 1727
- Release Information | 1727

Syntax

```
destination-host pcrf-hostname;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Destination-Host AVP value used in the CCR-GX-I message.

Options

pcrf-hostname—(Optional) Value of the Destination-Host AVP to be used in the CCR-GX-I message.

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 16.2.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

destination-realm (JSRC)

IN THIS SECTION

- Syntax | 1728
- Hierarchy Level | 1728
- Description | 1728
- Options | 1729
- Required Privilege Level | 1729
- Release Information | 1729

Syntax

```
destination-realm realm
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Configure the realm in which the SAE host resides.

Options

realm—Realm in which the SAE host resides.

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 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Configuring the JSRC Partition | 1477

destination-realm (Gx-Plus)

IN THIS SECTION

- Syntax | 1729
- Hierarchy Level | 1730
- Description | 1730
- Options | 1730
- Required Privilege Level | 1730
- Release Information | 1730

Syntax

```
destination-realm realm;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
```

Description

Configure the realm in which the PCRF host resides.

Options

realm—Realm in which the PCRF host resides.

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 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus | 1404

Configuring the Gx-Plus Partition | 1405

destination-realm (OCS Partition)

IN THIS SECTION

- Syntax | 1731
- Hierarchy Level | 1731
- Description | 1731

Options | 1731
Required Privilege Level | 1731
Release Information | 1731

Syntax

```
destination-realm ocs-realm-name;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the Destination-Realm AVP value in all CCR-GY messages.

Options

ocs-realm-name—Name of the Destination-Real AVP value to be used in all CCR-GY messages.

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 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition | 1449
3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
Understanding Gy Interactions Between the Router and the OCS | 1431
Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

destination-realm (PCRF Partition)

IN THIS SECTION

- Syntax | 1732
- Hierarchy Level | 1732
- Description | 1732
- Options | 1732
- Required Privilege Level | 1733
- Release Information | 1733

Syntax

```
destination-realm pcrf-realm-name;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Destination-Realm AVP value in all CCR-GX messages.

Options

pcrf-realmname—Value of the Destination-Real AVP to be used in all CCR-GX messages.

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 16.2.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

dhcp-attributes (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1733
- Hierarchy Level | 1734
- Description | 1734
- Options | 1735
- Required Privilege Level | 1739
- Release Information | 1739

Syntax

```
dhcp-attributes {
    boot-file filename;
```

```
    boot-server (address | hostname);
    dns-server [ ipv6-address ];
    domain-name domain-name;
    exclude-prefix-len exclude-prefix-length;
    grace-period seconds;
    maximum-lease-time seconds;
    name-server [ server-list ];
    netbios-node-type node-type;
    option {
        [ (id-number option-type option-value)
        (id-number array option-type option-value) ];
    }
    option-match {
        option-82 {
            circuit-id value range named-range;
            remote-id value range named-range;
        }
    }
    preferred-lifetime seconds;
    router [ router-address ];
    server-identifier ip4-address;
    sip-server-address [ ipv6-address ];
    sip-server-domain-name domain-name;
    t1-percentage percentage;
    t1-renewal-time;
    t2-percentage percentage;
    t2-rebinding-time;
    tftp-server address;
    valid-lifetime seconds;
    wins-server [ servers ];
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family family]
```

Description

Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

Options

boot-file

Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup. This configuration is equivalent to DHCP option 67.

- **Values:** *filename*—Location of the boot file on the boot server. The filename can include a pathname.

boot-server

Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup. This configuration is equivalent to DHCP option 66.

- Values:
  - *address*—IPv4 address of a boot server.
  - *hostname*—Fully qualified hostname of a boot server.

dns-server

Specify a DNS server to which clients can send DNS queries. This is equivalent to DHCPv6 option 23. To specify multiple DNS servers, add multiple `dns-server` statements in order of preference.

- **Values:** *ipv6-address*—IPv6 address of a DNS server.

domain-name

Configure the name of the domain in which clients search 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.

- **Values:** *domain-name*—Name of the domain.

exclude-prefix-len *exclude-prefix-length*

Specify the length of the IPv6 prefix to be excluded from the delegated prefix. Range: 1 through 128.

grace-period

Configure the amount of time that the client retains the address lease after the lease expires. The address cannot be reassigned to another client during the grace period.

- **Values:** *seconds*—Number of seconds the lease is retained.
- **Range:** 0 through 4,294,967,295 seconds.
- **Default:** 0 (no grace period).

maximum-lease-time

Specify the maximum length of time, in seconds, that the lease is held for a client if the client does not renew the lease. This is equivalent to DHCP option 51. The `maximum-lease-`

`time` is mutually exclusive with both the `preferred-lifetime` and the `valid-lifetime`, and cannot be configured with either timer.

- **Values:** *seconds*—Maximum number of seconds the lease can be held.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

name-server Configure one or more Domain Name System (DNS) name servers available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

- **Values:** *server-names*—IP addresses of the domain name servers, listed in order of preference.

netbios-node-type Specify the NetBIOS node type. This is equivalent to DHCP option 46.

- **Values:** *node-type*—One of the following node types:
  - b-node—Broadcast node.
  - h-node—Hybrid node.
  - m-node—Mixed node.
  - p-node—Peer-to-peer node.

option Specify user-defined options that are added to client packets. Starting in Junos OS Release 13.3, the `hex-string` option type was introduced.

- Values:
  - *array*—An option can include an array of option types.
  - *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
  - *option-type*—Any of the following types: byte, byte-stream, flag, hex-string, integer, ip-address, short, string, unsigned-integer, or unsigned-short.
  - *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an On or Off value for a flag type).

preferred-lifetime Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix active. When the lifetime expires, the address is deprecated. 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.

- **Values:** *seconds*—Number of seconds that the IPv6 prefix is active.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

router Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

- **Values:** *router-address*—IP address of one or more routers.

server-identifier Specify the IP address that is used as the source address the DHCP server includes in IP packets when communicating with clients. The address is included in the DHCP packet in option 54.

- **Values:** *ipv4-address*—IP address.

sip-server-address Specify a SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 22. To specify multiple servers, add multiple `sip-server-address` statements in order of preference.

- **Values:** *ipv6-address*—IPv6 address of a SIP outbound proxy server.

sip-server-domain-name Configure the domain name of the SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 21.

- **Values:** *domain-name*—Name of the domain.

t1-percentage Specify a percentage of the `preferred-lifetime` value. After this percentage of the `preferred-lifetime` value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from the originating DHCP local server. The `t1-percentage` is also referred to as the renewal time. The `t1-percentage` value must be less than the `t2-percentage` value. DHCPv4 server support was added in Junos OS Release 17.2.

- **Values:** *percentage*—Percentage of the `preferred-lifetime` value.
- **Range:** 0 through 100.
- **Default:** If the `t1-percentage` value is not configured, the default is based on the `preferred-lifetime` value:
  - If the `preferred-lifetime` value is finite, the default is 50 percent of the `preferred-lifetime` value.
  - If the `preferred-lifetime` value is infinite, the default is also infinite.

t1-renewal-time

Specify the time (T1) at which the DHCPv4 or DHCPv6 client requests an extension (renewal) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the `t1-percentage` statement.

- **Values:** *seconds*—Number of seconds.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 50 percent of the lease duration (`preferred-lifetime`).

t2-percentage

Specify a percentage of the `preferred-lifetime` value. After this percentage of the `preferred-lifetime` value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from any available DHCPv4 or DHCPv6 server. The `t2-percentage` is also referred to as the rebinding time. The `t2-percentage` value must be greater than the `t1-percentage` value. DHCPv4 server support was added in Junos OS Release 17.2.

- **Values:** *percentage*—Percentage of the `preferred-lifetime` value.
- **Range:** 0 through 100.
- **Default:** Default: If the `t2-percentage` value is not configured, the default is based on the `preferred-lifetime` value:
  - If the `preferred-lifetime` value is finite, the default is 80 percent of the `preferred-lifetime` value.
  - When the `preferred-lifetime` value is infinite, the default is also infinite.

t2-rebinding-time

Specify the time (T2) at which the DHCPv4 or DHCPv6 client attempts to contact any DHCP server to request an extension (rebinding) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the `t2-percentage` statement.

- **Values:** *seconds*—Number of seconds.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** The default value depends on the client:
  - (DHCPv4 clients) 87.5 percent of the lease duration (`preferred-lifetime`).
  - (DHCPv6 clients) 80 percent of the lease duration (`preferred-lifetime`).

tftp-server

Specify the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file. This is equivalent to DHCP option 150.

- **Values:** *ip-address*—IP address of the TFTP server.

valid-lifetime Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix valid. When the lifetime expires, the address becomes invalid. 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.

- **Values:** *seconds*—Number of seconds that the IPv6 prefix is valid.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

wins-server Specify one or more NetBIOS name servers (NBNS) that the client uses to resolve NetBIOS names. This is equivalent to DHCP option 44.

- **Values:** *ipv4-address*—IP address of each NetBIOS name server. Add them to the configuration in order of preference.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`exclude-prefix-len` statement introduced in Junos OS Release 17.3 for MX Series.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | **1138**

DHCP Client Attribute and Address Assignment | **766**

DHCP Lease Times for IP Addresses | **780**

dhcp-local-server

IN THIS SECTION

- Syntax | 1740
- Hierarchy Level | 1751
- Description | 1751
- Required Privilege Level | 1752
- Release Information | 1752

Syntax

```
dhcp-local-server {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
```

```
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dhcpv6 {
        access-profile profile-name;
        allow-active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
        }
        allow-bulk-leasequery {
            max-connections number-of-connections;
            max-empty-replies number-of-replies;
            restricted-requestor;
            timeout seconds;
        }
        allow-leasequery {
            restricted-requestor;
        }
        authentication {
            ...
        }
        duplicate-clients incoming-interface;
        group group-name {
            access-profile profile-name;
            authentication {
                ...
            }
            interface interface-name {
                access-profile profile-name;
                exclude;
                overrides {
                    asymmetric-lease-time seconds;
                    asymmetric-prefix-lease-time seconds;
                    delay-advertise {
                        based-on (option-15 | option-16 | option-18 | option-37) {
                            equals {
                                ascii ascii-string;
```

```
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
```

```
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    route-suppression;
    server-duid-type type;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
```

```
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
```

```
                }
                delay-time seconds;
            }
            delegated-pool;
            dual-stack dual-stack-group-name;
            include-option-82 {
                forcerenew;
                nak;
            }
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
        }
        reconfigure {
            attempts attempt-count;
            clear-on-terminate;
            strict;
            support-option-pd-exclude;
            timeout timeout-value;
            token token-value;
             trigger {
                radius-disconnect;
            }
        }
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        requested-ip-network-match subnet-mask;
        route-suppression;
        server-duid-type type;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group name {
        access-profile access-profile;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
```

```
                interface-description  (device-interface | logical-interface);
                interface-name ;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        on-demand-address-allocation;
        protocol-primary (inet | inet6);
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        service-profile service-profile;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    duplicate-clients-in-subnet  (incoming-interface | option-82);
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-profile-
name>;
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        authentication {
            ...
```

```
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-
profile-name>;
        interface interface-name {
            exclude;
            overrides {
                asymmetric-lease-time seconds;
                client-discover-match (option60-and-option82 | incoming-interface);
                delay-offer {
                    based-on (option-60 | option-77 | option-82) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                include-option-82 {
                    forcerenew;
                    nak;
                }
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
```

```
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            asymmetric-lease-time seconds;
            client-discover-match  (option60-and-option82 | incoming-interface);
            delay-offer {
                based-on (option-60 | option-77 | option-82) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                }
                delay-time seconds;
            }
```

```
            include-option-82 {
                forcerenew;
                nak;
            }
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
        }
        requested-ip-network-match subnet-mask
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    on-demand-address-allocation;
    overrides {
```

```
        asymmetric-lease-time seconds;
        client-discover-match  <option60-and-option82 | incoming-interface>;
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
    }
    pool-match-order {
        external-authority;
        ip-address-first;
        option-82;
    }
    protocol-primary;
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
```

```
    requested-ip-network-match subnet-mask;
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services],
[edit logical-systems logical-system-name system services],
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch to enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and, therefore, is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the `[edit forwarding-options helpers]` hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The `dhcpv6` stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.

NOTE: When you configure the `dhcp-local-server` statement at the routing instance hierarchy level, you must use a routing instance type of `virtual-router`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP | 691

DHCPv6 Local Server Overview | 910

dhcp-relay

IN THIS SECTION

- Syntax | 1752
- Hierarchy Level | 1767
- Description | 1767
- Required Privilege Level | 1768
- Release Information | 1768

Syntax

```
dhcp-relay {
    access-profile profile-name;
    active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
```

```
    }
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
stale-timer             vlan-tags;
        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    dhcpv6 {
        access-profile profile-name;
        active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
            topology-discovery;
        }
        active-server-group server-group-name;
        }
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
```

```
                mac-address mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        bulk-leasequery {
            attempts number-of-attempts;
            timeout seconds;
            trigger automatic;
        }
        duplicate-clients incoming-interface;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only-replies;
        }
        forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-
interfaces);
        group group-name {
            access-profile profile-name;
            active-server-group server-group-name;
            authentication {
                password password-string;
                username-include {
                    circuit-type;
                    client-id;
                    delimiter delimiter-character;
                    domain-name domain-name-string;
                    interface-description  (device-interface | logical-interface);
                    interface-name interface-name;
                    logical-system-name;
                    mac-address mac-address;
                    relay-agent-interface-id;
                    relay-agent-remote-id;
```

```
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    interface interface-name {
        access-profile profile-name;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        exclude;
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        trace;
        upto upto-interface-name;
    }
    }
    lease-time-validation {
        lease-time-threshold seconds;
```

```
            violation-action action;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
```

```
                keep-incoming-interface-id ;
                no-vlan-interface-name;
                prefix prefix;
                use-interface-description (logical | device);
                use-option-82 <strict>;
                use-vlan-id;
            }
            relay-agent-remote-id {
                include-irb-and-l2;
                keep-incoming-interface-id ;
                no-vlan-interface-name;
                prefix prefix;
                use-interface-description (logical | device);
                use-option-82 <strict>;
                use-vlan-id;
            }
            relay-option {
                option-number option-number;
                default-action {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
                equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
                starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
            }
            remote-id-mismatch disconnect;
            route-suppression;
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        }
        leasequery {
            attempts number-of-attempts;
            timeout seconds;
        }
```

```
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
        route-suppression;
        service-profile dynamic-profile-name;
    }
}
no-snoop;
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
```

```
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-remote-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        relay-option-vendor-specific{
            host-name;
            location;
        remote-id-mismatch disconnect;
```

```
        route-suppression;
        server-group {
            server-group-name {
                server-ip-address;
            }
        }
        server-response-time seconds;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group dual-stack-group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
```

```
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        protocol-primary (inet | inet6);
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-remote-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    duplicate-clients-in-subnet (incoming-interface | option-82):
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    forward-only-replies;
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
```

```
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 [circuit-id] [remote-id];
                routing-instance-name;
                user-prefix user-prefix-string;
                }
            vlan-tags;
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
```

```
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode (automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                }
            }
            overrides {
                allow-no-end-option;
                allow-snooped-clients;
                always-write-giaddr;
                always-write-option-82;
                asymmetric-lease-time seconds;
                client-discover-match <option60-and-option82 | incoming-interface>;
                delay-authentication;
                delete-binding-on-renegotiation;
                disable-relay;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                layer2-unicast-replies;
                no-allow-snooped-clients;
                no-bind-on-request;
                proxy-mode;
                relay-source
                replace-ip-source-with;
                send-release-on-delete;
                trust-option-82;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        overrides {
            allow-no-end-option
            allow-snooped-clients;
            always-write-giaddr;
            always-write-option-82;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match (option60-and-option82 | incoming-interface);
            delay-authentication;
```

```
            delete-binding-on-renegotiation;
            disable-relay;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            layer2-unicast-replies;
            no-allow-snooped-clients;
            no-bind-on-request;
            proxy-mode;
            relay-source
            replace-ip-source-with;
            send-release-on-delete;
            trust-option-82;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group group-name;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                local-server-group local-server-group;
                relay-server-group relay-server-group;
            }
        }
        relay-option-82 {
            circuit-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            remote-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            server-id-override
        }
```

```
        remote-id-mismatch disconnect;
        route-suppression:
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    lease-time-validation {
        lease-time-threshold seconds;
        violation-action action;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    no-snoop;
    overrides {
        allow-no-end-option
        allow-snooped-clients;
        always-write-giaddr;
```

```
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match (option60-and-option82 | incoming-interface);
        delay-authentication;
        delete-binding-on-renegotiation;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            relay-server-group group-name;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        remote-id {
```

```
            prefix prefix;
            use-interface-description (logical | device);
        }
        server-id-override
    }
    }
    remote-id-mismatch disconnect;
    route-suppression:
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the `dhcp-relay` and `dhcpv6` statements are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

- Extended DHCP Relay Agent Overview | 695
- DHCPv6 Relay Agent Overview | 916
- DHCP Relay Proxy Overview | 698
- Specifying Authentication Support | 833

dhcp-service

IN THIS SECTION

- Syntax | 1769
- Hierarchy Level | 1770
- Description | 1770
- Required Privilege Level | 1770
- Release Information | 1770

Syntax

```
dhcp-service {
    accept-max-tcp-connections max-tcp-connections;
    dhcp-snooping-file(local_pathname | remote_URL) {
        write-interval interval;
    }
    dhcpv6-snooping-file {
        location;
        write-interval seconds;
    }
    (disable | enable);
    interface-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;
    }
    log {
        session {
            client;
            all;
            dhcpv6 {
                client;
                server;
                relay;
                dynamic-server;
                all;
            }
            server;
            relay;
        }
    }
    ltv-syslog-interval seconds;
    persistent-storage {
        backup-interval backup-interval;
        file-name;
    }
    request-max-tcp-connections max-tcp-connections;
    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;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can inprove performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Support for log option introduced in Junos OS Release 19.1R1 for SRX Series devices.

RELATED DOCUMENTATION

- Configuring Persistent Bindings in the DHCP or DHCPv6 (ELS)

dhcpv6 (DHCP Local Server)

IN THIS SECTION

- Syntax | 1771
- Hierarchy Level | 1777
- Description | 1777
- Required Privilege Level | 1777
- Release Information | 1777

Syntax

```
dhcpv6 {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            logical-system-name;
```

```
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    duplicate-clients incoming-interface;
    group group-name {
        access-profile profile-name;
        authentication {
        ...
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode(automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                }
            }
            overrides {
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
                delay-advertise {
```

```
                    based-on (option-15 | option-16 | option-18 | option-37) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                delete-binding-on-renegotiation;
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
```

```
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
```

```
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
```

```
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
        reconfigure {
            attempts attempt-count;
            clear-on-terminate;
            strict;
            timeout timeout-value;
            token token-value;
             trigger {
                radius-disconnect;
            }
        }
    }
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        support-option-pd-exclude;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
    requested-ip-network-match subnet-mask;
    route-suppression;
    server-duid-type type;
    service-profile dynamic-profile-name;
```

```
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure DHCPv6 local server options on the router or switch to enable the router or switch to function as a server for the DHCP protocol for IPv6. The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

DHCPv6 Local Server Overview | 910

dhcpv6 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1778
- Hierarchy Level | 1785
- Description | 1785
- Required Privilege Level | 1785
- Release Information | 1785

Syntax

```
dhcpv6 {
    access-profile profile-name;
    active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
    }
    active-server-group server-group-name;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
```

```
            user-prefix user-prefix-string;
            vlan-tags;

        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
        trigger automatic;
    }
    duplicate-clients incoming-interface;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    forward-only-replies;
    }
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
```

```
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            dynamic-profile profile-name {
                aggregate-clients (merge | replace);
                use-primary primary-profile-name;
            }
            exclude;
            overrides {
                allow-snooped-clients;
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
                delay-authentication;
                delete-binding-on-renegotiation;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                no-allow-snooped-clients;
                no-bind-on-request;
                relay-source interface-name;
                send-release-on-delete;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        }
        lease-time-validation {
            lease-time-threshold seconds;
            violation-action action;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
```

```
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
```

```
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        remote-id-mismatch disconnect;
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    lease-time-validation {
        lease-time-threshold seconds;
        violation-action action;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
```

```
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
            route-suppression;
            service-profile dynamic-profile-name;
        }
    }
    no-snoop;
    overrides {
        allow-snooped-clients;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-authentication;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        relay-source interface-name;
        send-release-on-delete;
    }
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
```

```
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
    }
    relay-option-vendor-specific{
        host-name;
        location;
    remote-id-mismatch disconnect;
    route-suppression;
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
```

```
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Configure DHCPv6 relay options on the router or switch and enable the router or switch to function as a DHCPv6 relay agent. A DHCPv6 relay agent forwards DHCPv6 request and reply packets between a DHCPv6 client and a DHCPv6 server.

The DHCPv6 relay agent server is fully compatible with the extended DHCP local server and DHCP relay agent. However, the options configured with the `dhcpv6` statement are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, the DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot be enabled on the router or switch at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for `forward-snooped-clients` introduced in Junos OS Release 15.1X53-D56 for EX Series switches and Junos OS Release 17.1R1.

RELATED DOCUMENTATION

dhcp-relay | 1752

DHCPv6 Relay Agent Overview | 916

Specifying Authentication Support | 833

diameter

IN THIS SECTION

- Syntax | 1786
- Hierarchy Level | 1787
- Description | 1787
- Options | 1787
- Required Privilege Level | 1788
- Release Information | 1788

Syntax

```
diameter {
    network-element element-name {
        dne-origin realm realm-name <host hostname>;
        forwarding {
            route dne-route-name {
                destination realm realm-name <host hostname>;
                function function-name <partition partition-name>;
                metric route-metric;
            }
        }
        function function-name;
        peer peer-name {
            priority priority-number;
        }
    }
    origin realm realm-name host hostname;
```

```
    peer peer-name {
        address ip-address;
        connect-actively {
            port port-number;
            transport transport-name;
        }
        logical-system logical-system-name  <routing-instance routing-instance-name> ;
        peer-origin realm realm-name host hostname;
        routing-instance routing-instance-name;
    }
    product-name name;
    transport transport-name {
        address;
        logical-system logical-system-name <routing-instance routing-instance-name>;
        routing-instance routing-instance-name;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure the Diameter base protocol for subscriber management.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

product-name *name* Name of a product to advertise in the Capability Exchange Request and Capability-Exchange Answer messages to a Diameter peer. The name is included in the Diameter AVP Product-Name (269). If you change the product-name, the MX series router acting as a BNG disconnects all existing connections to Diameter peers and reconnects to those peers using the new product name.

- **Default:** JUNOS

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 9.6.

`product-name` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Configuring Diameter | 1373

diameter-instance (JSRC)

IN THIS SECTION

- Syntax | 1788
- Hierarchy Level | 1789
- Description | 1789
- Options | 1789
- Required Privilege Level | 1789
- Release Information | 1789

Syntax

```
diameter-instance instance-name
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Specify the Diameter instance associated with the JSRC partition.

Options

instance-name—Name of the Diameter instance. Currently, only `master` is supported.

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 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Configuring the JSRC Partition | 1477

diameter-instance (Diameter Applications)

IN THIS SECTION

- Syntax | 1790
- Hierarchy Level | 1790
- Description | 1790

- Options | 1790
- Required Privilege Level | 1790
- Release Information | 1790

Syntax

```
diameter-instance instance-name;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
[edit access ocs partition partition-name],
[edit access pcrf partition partition-name]
```

Description

Specify the Diameter instance associated with the Gx-Plus, OCS, or PCRF partition.

Options

instance-name—Name of the Diameter instance. Currently, only `master` is supported.

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 11.2.

Support at the `[edit access ocs partition partition-name]` and `[edit access pcrf partition partition-name]` hierarchy levels introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- Configuring Gx-Plus | 1404
- Configuring the Gx-Plus Partition | 1405
- Configuring the OCS Partition | 1449
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Gy Interactions Between the Router and the OCS | 1431

dictionary

IN THIS SECTION

- Syntax | 1791
- Hierarchy Level | 1791
- Description | 1791
- Options | 1792
- Required Privilege Level | 1792
- Release Information | 1792

Syntax

```
dictionary dictionary-path;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Configure an XML-based dictionary file. The dictionary path is the complete path to the dictionary file and includes the dictionary filename. Extensible Subscriber Services Manager acts on the extensible-subscriber-service request on the basis of the services configured in the dictionary file. This configuration is mandatory.

Options

dictionary-path Path to the dictionary file. The complete path including the filename must not be more than 127 characters.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- *show extensible-subscriber-services dictionary*
- *show extensible-subscriber-services dictionary attributes*
- request services extensible-subscriber-services reload-dictionary | **2646**
- *show extensible-subscriber-services dictionary services*
- *Understanding the Dictionary File*

disable

IN THIS SECTION

- Syntax | **1793**
- Hierarchy Level | **1793**
- Description | **1793**
- Options | **1793**
- Required Privilege Level | **1793**
- Release Information | **1793**

Syntax

```
disable service-name;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Disable the service name of the subscriber profile.

Options

service-name—Name of the disabled service.

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 11.4.

disable (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | **1794**
- Hierarchy Level | **1794**
- Description | **1794**
- Default | **1794**

Required Privilege Level | 1794
Release Information | 1794

Syntax

```
disable
```

Hierarchy Level

```
[edit system processes extensible-subscriber-services]
```

Description

Disable the Extensible Subscriber Services Manager process when there is an extensible-subscriber-services configuration and the user wants to stop the process. The process is disabled by default when there is no configuration under the `[edit system services extensible-subscriber-services]` hierarchy level.

Default

The process is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

restart extensible-subscriber-services | 2667

disable-relay

IN THIS SECTION

- Syntax | 1795
- Hierarchy Level | 1795
- Description | 1795
- Required Privilege Level | 1796
- Release Information | 1796

Syntax

```
disable-relay;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Disable DHCP relay on specific interfaces in a group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

dne-origin (Diameter Network Element)

IN THIS SECTION

- Syntax | 1796
- Hierarchy Level | 1796
- Description | 1797
- Options | 1797
- Required Privilege Level | 1797
- Release Information | 1797

Syntax

```
dne-origin realm realm-name <host hostname>;
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Specify values of Origin-Realm-AVP and Origin-Host-AVP used in messages sent for the specified network element by the Diameter instance.

NOTE: Only the realm is mandatory for the DNE origin.

Options

host *hostname* (Optional) Name of the message origin host that is supplied as the value of the Origin-Host AVP for Diameter messages associated with the network element.

realm *realm-name* Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for Diameter messages associated with the network element.

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 16.2.

RELATED DOCUMENTATION

Configuring Diameter Network Elements | 1377

Configuring Diameter | 1373

dns-server-address (Dynamic Profiles)

IN THIS SECTION

- Syntax | 1798
- Hierarchy Level | 1798
- Description | 1798
- Options | 1799
- Required Privilege Level | 1799
- Release Information | 1799

Syntax

```
dns-server-address [dns-server-address | $junos-ipv6-dns-server-address] {
    lifetime seconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-
name]
```

Description

Specify the address of the DNS server that is used to resolve IPv6 DNS names. You can use RADIUS to provide the address dynamically in the $junos-ipv6-dns-server variable within Access-Accept messages, or you can statically configure up to three IPv6 addresses.

You can also specify the maximum time in seconds for which the DNS server address remains valid. The device can use the recursive DNS server address for DNS name resolution until the time specified expires.

Options

$junos-ipv6-dns-server-address — Dynamically receive the address of the DNS server from RADIUS in Access-Accept messages.

dns-server-address — IPv6 address of the DNS server. You can configure up to three DNS server addresses.

lifetime ***seconds*** — Maximum time for which the recursive DNS server address remains valid.

- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 1800 seconds
- **Values:** 0 indicates that the advertised recursive DNS server address is no longer valid and that this recursive DNS server address entry can be deleted.

  A value of 4,294,967,295 seconds indicates an infinite lifetime and a persistent entry in the device for this recursive DNS server address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

DNS Resolver for IPv6 DNS Overview | **1170**

Configuring a DNS Server Address for IPv6 Hosts | **1170**

domain (Domain Map)

IN THIS SECTION

- Syntax | 1800
- Hierarchy Level | 1801
- Description | 1801
- Required Privilege Level | 1801
- Release Information | 1801

Syntax

```
domain {
    delimiter [delimiter-character];
    map domain-map-name {
        aaa-logical-system logical-system-name {
            aaa-routing-instance routing-instance-name;
        }
        aaa-routing-instance routing-instance-name;
        access-profile profile-name;
        address-pool pool-name;
        dynamic-profile profile-name;

        strip-domain;
        target-logical-system logical-system-name {
            target-routing-instance routing-instance-name;
        }
        target-routing-instance routing-instance-name;
        tunnel-profile profile-name;
    }
    parse-direction (left-to-right | right-to-left);
    parse-order (domain-first | realm-first);
    realm-delimiter [delimiter-character];
    realm-parse-direction  (left-to-right | right-to-left);
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a domain map, which is used to map access options and session parameters for subscriber sessions.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring a Domain Map | 659

domain-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 1802
- Hierarchy Level | 1802
- Description | 1803
- Options | 1803
- Required Privilege Level | 1803

Release Information | 1803

Syntax

```
domain-name domain-name-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
```

```
group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

domain-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1804
- Hierarchy Level | 1804
- Description | 1805
- Options | 1805
- Required Privilege Level | 1805
- Release Information | 1805

Syntax

```
domain-name domain-name-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 authentication username-include],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

domain-name (Static Subscribers)

IN THIS SECTION

- Syntax | 1806
- Hierarchy Level | 1806
- Description | 1807
- Options | 1807
- Required Privilege Level | 1807
- Release Information | 1807

Syntax

```
domain-name domain-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify the domain name that is included at the end of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version.

Options

domain-name—Domain name that ends the username created for all static subscribers. The username is also sent to RADIUS in the Access-Request message. The string can include the following characters: a through z, A through Z, 0 through 9, "-", or ".".

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487

Configuring the Static Subscriber Global Username | 1492

Configuring the Static Subscriber Group Username | 1497

domain-name-server (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 1808
- Hierarchy Level | 1808
- Description | 1808

- Options | 1808
- Required Privilege Level | 1808
- Release Information | 1809

Syntax

```
domain-name-server dns-address;
```

Hierarchy Level

```
[edit access];
[edit access profile]
```

Description

Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

NOTE: A DNS name server address configured with this statement is less preferred than one configured with the "domain-name-server-inet" on page 1809 statement. That is, the server with the address configured with the "domain-name-server-inet" on page 1809 takes precedence over a server configured with this statement.

Options

dns-address	IPv4 address of the DNS name server.

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 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management | 1167

DNS Name Server Address Overview | 1165

domain-name-server-inet (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 1809
- Hierarchy Level | 1809
- Description | 1810
- Options | 1810
- Required Privilege Level | 1810
- Release Information | 1810

Syntax

```
domain-name-server-inet dns-address;
```

Hierarchy Level

```
[edit access],
[edit access profile]
```

Description

Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

NOTE: A DNS name server address configured with this statement is higher in preference than one configured with the "domain-name-server" on page 1807 statement. That is, the server with the address configured with the "domain-name-server-inet" on page 1809 takes precedence over a server configured with the "domain-name-server" on page 1807 statement.

Options

dns-address	IPv4 address of the DNS name server.

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 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management | 1167

DNS Name Server Address Overview | 1165

domain-name-server-inet6 (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 1811
- Hierarchy Level | 1811
- Description | 1811
- Options | 1811
- Required Privilege Level | 1811
- Release Information | 1812

Syntax

```
domain-name-server-inet6 dns-address;
```

Hierarchy Level

```
[edit access],
[edit access profile]
```

Description

Configure an IPv6 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

Options

dns-address	IPv6 address of the DNS name server.

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 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management | 1167

DNS Name Server Address Overview | 1165

downstream-rate (Traffic Shaping)

IN THIS SECTION

- Syntax | 1812
- Hierarchy Level | 1812
- Description | 1813
- Options | 1813
- Required Privilege Level | 1813
- Release Information | 1813

Syntax

```
downstream-rate rate;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit advisory-options],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name advisory-options],
```

```
[edit interfaces demux0 unit logical-unit-number advisory-options],
[edit interfaces interface-name unit logical-unit-number advisory-options]
```

Description

Specify a recommended shaping rate to be applied to downstream traffic on an interface.

For ANCP interfaces, this configured rate is used as the default value for the Juniper VSA Downstream-Calculated-Qos-Rate (26–141) when the router has not received and processed the attributes from the access node.

For L2TP, the rate is configured on an underlying PPPoE logical interface for a subscriber on an MX Series router acting as a LAC. When the subscriber is tunneled, this rate, referred to as speed for L2TP, is sent to the LNS in the ICCN message as AVP 24.

Options

rate—Traffic rate in bits per second.

- **Range:** 1000 through 4,294,967,295 bits per second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | **1304**

Configuring the ANCP Agent | **1253**

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

draining (Diameter Applications)

IN THIS SECTION

- Syntax | 1814
- Hierarchy Level | 1814
- Description | 1814
- Required Privilege Level | 1815
- Release Information | 1815

Syntax

```
draining;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) or Online Charging System (OCS) partition to the draining state to make substantial configuration changes quickly. To log out all subscribers quickly and immediately, issue the `clear network-access pcrf subscribers` command.

After you set draining for either of the PCRF or OCS partitions, any new subscriber logins are denied, and the time limit you set in the `draining-response-timeout` statement is used instead of the `logout-response-timeout` time limit.

After the partition has drained, you must issue the `clear network-access pcrf` or `clear network-access ocs statistics` command to clear out the respective partition's subscribers and wait for all of the subscribers to log out before you can make configuration changes and resume normal operations.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

draining-response-timeout (Diameter Applications)

IN THIS SECTION

- Syntax | 1816
- Hierarchy Level | 1816
- Description | 1816
- Options | 1816
- Required Privilege Level | 1816
- Release Information | 1816

Syntax

```
draining-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name]
```

Description

Configure the amount of time in seconds before a PCRF or an Online Charging System OCS partition responds and begins to drain. Configuring this statement is optional if you set `draining` statement for either of the Policy and Charging Rule Function (PCRF) or Online Charging System (OCS) partitions.

After you set draining for either of the PCRF or OCS partitions, any new subscriber logins are denied, and the time limit you set in this statement is used instead of the `logout-response-timeout` time limit.

Options

seconds Number of seconds to wait before the partition begins to drain.

- **Default:** 400
- **Range:** 0 through 86,400 seconds (24 hours)

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

drop (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 1817
- Hierarchy Level | 1817
- Description | 1818
- Required Privilege Level | 1818
- Release Information | 1818

Syntax

```
drop;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with)],
[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 ...]
```

Description

Drop (discard) specified DHCP client packets when you use DHCP relay agent selective processing. You can configure the drop operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

dsl (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 1819
- Hierarchy Level | 1820
- Description | 1820
- Options | 1821
- Required Privilege Level | 1823
- Release Information | 1823

Syntax

```
dsl {
    adsl {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    adsl2 {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    adsl2-plus {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    gfast {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    gfast-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    other {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    sdsl {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    sdsl-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    type tlv-value {
        overhead-adjust percentage;
```

```
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2 {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-annex-q {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-annex-q-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-bonded {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Configure adjustments to actual DSL data rates as follows:

- Multiply the actual data rate by a percentage for only downstream rates or for both upstream an downstream rates

- Adjust the encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes

The actual (unadjusted) downstream and upstream data rates, DSL line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

If the DSL-Type TLV (0x91) is not received in either the ANCP Port Status message or PPPoE-IA tags, the default adjustment leaves the rates and bytes unchanged.

Adjustments are applied to all subscribers using access lines of the specific DSL line type. Depending on the value, it may be reported to AAA, CoS, or both:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate. A change in value triggers an immediate Interim-Accounting message to the RADIUS server if you have configured the "ancp-speed-change-immediate-update" on page 1621 statement.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following AVPs and messages:

- AVP 24, Tx Connect Speed (ICCN message)—For the initial total rate adjustment to Actual-Data-Rate-Downstream TLV (0x82).
- AVP 38, Rx Connect Speed (ICCN message)—For the initial total rate adjustments to Actual-Data-Rate-Upstream TLV (0x81).
- AVP 97, Connect Speed Update (CSUN message)—For subsequent changes to the initial rates reported in AVP 24 and AVP 38.

Options

adsl Sets attributes for ADSL access lines.

adsl2 Sets attributes for ADSL2 access lines.

adsl2-plus Sets attributes for ADSL2+ access lines.

gfast Sets attributes for G.fast high speed access lines connected to a PON tree infrastructure.

gfast-bonded Sets attributes for G.fast high speed bonded access lines connected to a PON tree infrastructure.

other Sets attributes for access lines of type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

overhead-adjust *percentage* Adjusts the actual downstream rates for all subscribers on the specified access line by multiplying the rate by the specified percentage. This adjustment accounts for the Layer 1 overhead for the DSL type. For the subscriber access line, the adjustment is made to the downstream rate reported in the Actual-Data-Rate-Downstream TLV (0x82). The overhead is adjusted as follows:

- When Agent-Remote-Id TLV (0x02) is different than Access-Aggregation-Circuit-ID-ASCII (0x03), the adjustment shapes the logical interface (residential subscribers) or the interface set (business subscribers), as determined by the CoS adjustment control profile.
- When Agent-Remote-Id TLV (0x02) equals Access-Aggregation-Circuit-ID-ASCII (0x03), the adjustment shapes the PON tree or bonded copper line (parent) interface set, as determined by the CoS adjustment control profile.
- **Range:** 80 through 100 percent
- **Default:** 100 percent

overhead-bytes *bytes* Adjusts the actual downstream cell overhead for all subscribers on the specified access line by adding or subtracting the specified number of bytes. The adjustment accounts for the traffic encapsulation overhead. The adjustment is made to the overhead reported in the Actual-Data-Rate-Downstream TLV (0x82). The overhead is adjusted as follows:

- Agent-Remote-Id TLV (0x02) is different than Access-Aggregation-Circuit-ID-ASCII (0x03)—The adjustment shapes the logical interface (residential subscribers) or the interface set (business subscribers), as determined by the CoS adjustment control profile.
- Agent-Remote-Id TLV (0x02) equals Access-Aggregation-Circuit-ID-ASCII (0x03)—The adjustment shapes the (parent) interface set for a PON tree (FTTB) or bonded copper line (CuTTB) as determined by the CoS adjustment control profile.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl Sets attributes for SDSL access lines.

sdsl-bonded Sets attributes for bonded SDSL access lines.

total-adjust *percentage* Adjusts the downstream and upstream data rates for all subscribers on an access line of the specified types by multiplying the rate by the specified percentage. This adjustment accounts for the total Layer 1 and encapsulation overhead for the DSL type. The adjustment is made to the Actual-Data-Rate-Downstream TLV (0x82) and Actual-Data-Rate-Upstream TLV (0x81).

- **Range:** 1 through 100 percent
- **Default:** 100 percent

type *tlv-value* Sets attributes for access lines by specifying the unsigned integer value of the DSL-Type TLV (0x91) to reference the DSL type. This option enables the `dsl` statement to be used for DSL access line types that might be introduced in the future.

- **Range:** 14 through 4294967295

vdsl Sets attributes for VDSL access lines.

vdsl2 Sets attributes for VDSL2 access lines.

vdsl2-annex-q Sets attributes for VDSL2 Annex Q access lines.

vdsl2-annex-q-bonded Sets attributes for bonded VDSL2 Annex Q access lines.

vdsl2-bonded Sets attributes for bonded VDSL access lines.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | 1308

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

dual-stack (DHCP Local Server Overrides)

IN THIS SECTION

- Syntax | 1824
- Hierarchy Level | 1824
- Description | 1825
- Options | 1825
- Required Privilege Level | 1825
- Release Information | 1825

Syntax

```
dual-stack dual-stack-group-name;
```

Hierarchy Level

```
[edit logical-systems name routing-instances routing-instance-name system services dhcp-local-server ...],
[edit logical-systems name system services dhcp-local-server ...],
[edit routing-instances name system services dhcp-local-server ...],
[edit system services dhcp-local-server group group-name interface interface-name overrides],
```

```
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server overrides]
```

Description

Assigns the specified dual-stack group to both legs (DHCP and DHCPv6) of the DHCP dual stack. The dual-stack group defines the common configuration settings for DHCP and DHCPv6 subscribers on both legs. These settings take precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

dual-stack-group-name	Name of the globally configured dual-stack group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview | 1004

Configuring Single-Session DHCP Dual-Stack Support | 1008

dual-stack (DHCP Relay Agent Overrides)

IN THIS SECTION

- Syntax | 1826
- Hierarchy Level | 1826

- Description | 1826
- Options | 1826
- Required Privilege Level | 1827
- Release Information | 1827

Syntax

```
dual-stack dual-stack-group-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Assigns the specified dual-stack group to both legs (DHCP and DHCPv6) of the DHCP dual stack. The dual-stack group defines the common configuration settings for DHCP and DHCPv6 subscribers on both legs. These settings take precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

dual-stack-group-name	Name of the globally configured dual-stack group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview | 1004

Configuring Single-Session DHCP Dual-Stack Support | 1008

dual-stack-group (DHCP Local Server)

IN THIS SECTION

- Syntax | 1827
- Hierarchy Level | 1829
- Description | 1829
- Options | 1829
- Required Privilege Level | 1829
- Release Information | 1829

Syntax

```
dual-stack-group name {
    access-profile access-profile;
    authentication {
        password password-string;
        username-include {
            circuit-type;
```

```
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    on-demand-address-allocation;
    protocol-primary (inet | inet6);
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    service-profile service-profile;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server],
[edit logical-systems name system services dhcp-local-server],
[edit routing-instances name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP local server dual-stack, and names the dual-stack group.

When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

name Name of the dual-stack group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- Single-Session DHCP Dual-Stack Overview | **1004**
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets
- Configuring RADIUS Reauthentication for DHCP Subscribers | **567**
- RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | **556**

dual-stack-group (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1830
- Hierarchy Level | 1831
- Description | 1832
- Options | 1832
- Required Privilege Level | 1832
- Release Information | 1832

Syntax

```
dual-stack-group name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
```

```
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    protocol-primary (inet | inet6);
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    service-profile dynamic-proflle-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP dual stack, and names the dual stack group.

The group is assigned to each leg of the DHCP dual-stack with the "dual-stack" on page 1825 statement in the `overrides` stanza. When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

name	Name of the dual-stack group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- Single-Session DHCP Dual-Stack Overview | **1004**
- Configuring Single-Session DHCP Dual-Stack Support | **1008**
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 1833
- Hierarchy Level | 1833
- Description | 1833
- Options | 1834
- Required Privilege Level | 1834
- Release Information | 1834

Syntax

```
dual-stack-interface-client-limit number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dual-stack-group
dual-stack-group-name ],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-
group dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit routing-instances name forwarding-options dhcp-relay dual-stack-group dual-stack-group-
name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit system services dhcp dhcp-local-server dual-stack-group dual-stack-group-name ],
```

Description

Limit the number of clients allowed on an interface.

NOTE: For dual-stack subscribers, always use this statement instead of the "interface-client-limit" on page 1958 (DHCP Relay Agent) or "interface-client-limit" on page 1955 (DHCP Local Server) statements.

Options

number Maximum number of dual-stack subscribers that can log in per interface.

- **Range:** 1 through 500,000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview | **1004**

dualstack-support (JSRC)

IN THIS SECTION

- Syntax | **1835**
- Hierarchy Level | **1835**
- Description | **1835**
- Default | **1835**
- Required Privilege Level | **1836**

Release Information | 1836

Syntax

```
dualstack-support;
```

Hierarchy Level

```
[edit jsrc]
```

Description

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. By default (and in all cases for releases earlier than Junos OS Release 18.1R1), 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.

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 Framed-IP-Address (AVP 8) and Framed-IP-Netmask (AVP 9). For IPv6, that includes Framed-IPv6-Address (AVP 168), Framed-IPv6-Prefix (AVP 97), Delegated-IPv6-Prefix (AVP 123), Juniper-IPv6-Ndra-Prefix (AVP 2200), and 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.

Default

Disabled.

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 18.1R1.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

duplication (Access Profile)

IN THIS SECTION

- Syntax | 1836
- Hierarchy Level | 1836
- Description | 1837
- Default | 1837
- Required Privilege Level | 1837
- Release Information | 1837

Syntax

```
duplication;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send accounting reports to both the RADIUS accounting server configured in the access profile for the wholesaler and the RADIUS accounting server configured in the access profile for the retailer.

Default

The router sends accounting reports to the accounting servers that are in the context in which the subscriber is authenticated.

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 11.4.

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access | 550

Understanding RADIUS Accounting Duplicate Reporting | 578

duplication-filter (Access Profile)

IN THIS SECTION

- Syntax | 1838
- Hierarchy Level | 1838
- Description | 1838
- Options | 1838
- Required Privilege Level | 1838

Release Information | 1839

Syntax

```
duplication-filter (interim-duplicated | interim-original) <exclude-attributes>;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure which accounting servers receive accounting messages when RADIUS duplication accounting is configured in the access profile.

Options

exclude-attributes (Optional) Filter RADIUS attributes from duplicated accounting interim messages based on the `exclude` statement configuration in the corresponding duplication access profile. You can configure `exclude-attributes` alone, or with `interim-duplicated` or `interim-original`.

interim-duplicated Do not send accounting interim messages to RADIUS accounting servers that are in a duplication context other than the subscriber's access profile.

interim-original Do not send accounting interim messages to RADIUS accounting servers that are in the subscriber's access profile.

NOTE: The `interim-duplicated` and `interim-original` filters are mutually exclusive.

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 16.1.

RELATED DOCUMENTATION

- Configuring Authentication and Accounting Parameters for Subscriber Access | 550
- Understanding RADIUS Accounting Duplicate Reporting | 578
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting | 580
- Configuring Per-Subscriber Session Accounting | 573

duplication-vrf (Duplicate Accounting)

IN THIS SECTION

- Syntax | 1839
- Hierarchy Level | 1840
- Description | 1840
- Required Privilege Level | 1840
- Release Information | 1840

Syntax

```
duplication-vrf {
    access-profile-name profile-name;
    vrf-name vrf-name;
}
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send duplicate accounting information to the RADIUS accounting servers defined in up to five access profiles all in the same nondefault VRF (LS:RI combination).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

Understanding RADIUS Accounting Duplicate Reporting | 578

Configuring Authentication and Accounting Parameters for Subscriber Access | 550

dynamic-profile (Demux)

IN THIS SECTION

- Syntax | 1841
- Hierarchy Level | 1841

- Description | 1841
- Options | 1841
- Required Privilege Level | 1842
- Release Information | 1842

Syntax

```
dynamic-profile profile-name {
    network ip-address {
        range name {
            low lower-limit;
            high upper-limit;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges]
```

Description

Assign a dynamic profile and specify address options for the demultiplexing (demux) interface options.

Options

profile-name	Name of the dynamic profile for the demultiplexing (demux) interface options.
network *ip-address*	Configure an IPv4 or IPv6 address for a dynamic profile for the demultiplexing (demux) interface options.
range <range-name>	Configure an IP name range used within an address-assignment pool for the demultiplexing (demux) interface options.

- `low` *lower-limit*—Lower limit of IPv4 or IPv6 address range.
- `high` *upper-limit*—Upper limit of IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | **1110**

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | **1111**

dynamic-profile (DHCP Local Server)

IN THIS SECTION

- Syntax | **1843**
- Hierarchy Level | **1843**
- Description | **1843**
- Options | **1843**
- Required Privilege Level | **1843**
- Release Information | **1844**

Syntax

```
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[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 ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Options `aggregate-clients` and `use-primary` introduced in Junos OS Release 9.3.

Support at the `[edit ... interface]` hierarchy levels introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
>
> Configuring a Default Subscriber Service | 765

dynamic-profile (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1844
- Hierarchy Level | 1845
- Description | 1845
- Options | 1845
- Required Privilege Level | 1845
- Release Information | 1845

Syntax

```
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.

M120 and M320 routers do not support DHCPv6.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 1752

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Grouping Interfaces with Common DHCP Configurations | 852

Configuring a Default Subscriber Service | 765

dynamic-profile (Domain Map)

IN THIS SECTION

- Syntax | 1846
- Hierarchy Level | 1846
- Description | 1846
- Options | 1847
- Required Privilege Level | 1847
- Release Information | 1847

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 | 664

dynamic-profile (Static Subscribers)

IN THIS SECTION

- Syntax | 1848
- Hierarchy Level | 1848
- Description | 1848
- Default | 1848
- Options | 1848
- Required Privilege Level | 1849
- Release Information | 1849

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]
```

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 | 1487
- Specifying the Static Subscriber Global Dynamic Profile | 1490
- Specifying the Static Subscriber Group Dynamic Profile | 1495

dynamic-profiles

IN THIS SECTION

- Syntax | 1850
- Hierarchy Level | 1861
- Description | 1861
- Options | 1861
- Required Privilege Level | 1862
- Release Information | 1862

Syntax

```
dynamic-profiles {
    profile-name {
        class-of-service {
            dynamic-class-of-service-options {
                vendor-specific-tags tag;
            }
            interfaces {
                interface-name ;
                }
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-
profile);
                    report-ingress-shaping-rate bps;
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                        }
                    }
                }
            }
            scheduler-maps {
                map-name {
                    forwarding-class class-name scheduler scheduler-name;
                }
            }
            schedulers {
                (scheduler-name) {
                    buffer-size (seconds | percent percentage | remainder | temporal
microseconds);
                    drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
protocol (any | non-tcp | tcp) drop-profile profile-name;
                    excess-priority (low | high | $junos-cos-scheduler-excess-priority);
                    excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
                    overhead-accounting (shaping-mode) <bytes (byte-value>;
                    priority priority-level;
```

```
                    shaping-rate (rate | predefined-variable);
                    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
                }
            }
            traffic-control-profiles profile-name {
                adjust-minimum rate;
                delay-buffer-rate (percent percentage | rate);
                excess-rate (percent percentage | proportion value | percent $junos-cos-excess-
rate);
                excess-rate-high (percent percentage | proportion value);
                excess-rate-low (percent percentage | proportion value);
                guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
                max-burst-size cells;
                overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
                peak-rate rate;
                scheduler-map map-name;
                shaping-rate (percent percentage | rate | predefined-variable) <burst-size
bytes>;
                shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
                sustained-rate rate;
            }
        }
        firewall {
            family family {
                fast-update-filter filter-name {
                    interface-specific;
                    match-order [match-order];
                    term term-name {
                        from {
                            match-conditions;
                        }
                        then {
                            action;
                            action-modifiers;
                        }
```

```
                    only-at-create;
                }
            }
            filter filter-name {
                enhanced-mode-override;
                instance-shared;
                interface-shared;
            interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                only-at-create;
            filter filter-name {
            interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
        hierarchical-policer uid {
            aggregate {
                if-exceeding {
                    bandwidth-limit-limit bps;
                    burst-size-limit bytes;
                }
                then {
                    policer-action;
                }
            }
            premium {
                if-exceeding {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
```

```
                then {
                    policer-action;
                }
            }
        }
        policer uid {
            filter-specific;
            if-exceeding {
                (bandwidth-limit bps | bandwidth-percent percentage);
                burst-size-limit bytes;
            }
            logical-bandwidth-policer;
            logical-interface-policer;
            physical-interface-policer;
            then {
                policer-action;
            }
        }
        three-color-policer uid {
            action {
                loss-priority high then discard;
            }
            logical-interface-policer;
            single-rate {
                (color-aware | color-blind);
                committed-burst-size bytes;
                committed-information-rate bps;
                excess-burst-size bytes;
            }
            two-rate {
                (color-aware | color-blind);
                committed-burst-size bytes;
                committed-information-rate bps;
                peak-burst-size bytes;
                peak-information-rate bps;
                }
            }
        }
    }
    interfaces interface-name {
        interface-set interface-set-name {
            interface interface-name {
                unit logical unit number {
```

```
                        advisory-options {
                            downstream-rate rate;
                            upstream-rate rate;
                        }
                    }
                }
            }
            unit logical-unit-number {
                actual-transit-statistics;
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                    line-identity {
                        include {
                            accept-no-ids;
                            circuit-id;
                            remote-id;

                        }
                        dynamic-profile profile-name;
                    }
                }
                encapsulation  (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-
vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap |
atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-
ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether |
ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
                family family {
                    address address;
                    filter {
                        adf {
                            counter;
                            input-precedence precedence;
                            not-mandatory;
                            output-precedence precedence;
                            rule rule-value;
                        }
                        input filter-name (
                            precedence precedence;
                            shared-name filter-shared-name;
                        }
```

```
                output filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
            }
            rpf-check {
                fail-filter filter-name;
                mode loose;
            }
            service {
                 input {
                     service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                input-vlan-map {
                    inner-tag-protocol-id tpid;
                    inner-vlan-id number;
                    (push | swap);
                    tag-protocol-id tpid;
                    vlan-id number;
                }
                output {
                     service-set service-set-name {
                        service-filter filter-name;
                    }
                }
                output-vlan-map {
                    inner-tag-protocol-id tpid;
                    inner-vlan-id number;
                    (pop | swap);
                    tag-protocol-id tpid;
                    vlan-id number;
                }
                pcef  pcef-profile-name {
                    activate rule-name | activate-all;
                }
            }
            unnumbered-address interface-name <preferred-source-address address>;
        }
        filter {
            input filter-name (
```

```
                shared-name filter-shared-name;
            }
            output filter-name {
                shared-name filter-shared-name;
            }
        }
        host-prefix-only;
        ppp-options {
            aaa-options aaa-options-name;
            authentication [ authentication-protocols ];
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
                local-name name;
            }
            ignore-magic-number-mismatch;
            initiate-ncp (dual-stack-passive | ipv6 | ip)
            ipcp-suggest-dns-option;
            mru size;
            mtu (size | use-lower-layer);
            on-demand-ip-address;
            pap;
            peer-ip-address-optional;
            local-authentication {
                password password;
                username-include {
                    circuit-id;
                    delimiter character;
                    domain-name name;
                    mac-address;
                    remote-id;
                }
            }
        }
        reassemble-packets;
        targeted-options {
            backup backup;
            group group;
            primary primary;
            weight ($junos-interface-target-weight | weight-value);
        }
        telemetry {
            subscriber-statistics;
            queue-statistics {
```

```
                    interface $junos-interface-name {
                        refresh rate;
                        queues queue set;
                    }
                    interface-set $junos-interface-set-name {
                        refresh rate;
                        queues queue set;
                    }
                }
            }
            vlan-id number;
            vlan-tags outer  [tpid].vlan-id [inner [tpid].vlan-id];
        }
    }
    interfaces {
        demux0 {...}
    }
    interfaces {
        pp0 {...}
    }
    policy-options {
        prefix-list uid {
            ip-addresses;
            dynamic-db;
        }
    }
    predefined-variable-defaults predefined-variable <variable-option> default-value;
    profile-type remote-device-service;
    protocols {
        igmp {
            interface interface-name {
                accounting;
                disable;
                group-limit limit;
                group-policy;
                group-threshold value;
                immediate-leave
                log-interval seconds;
                no-accounting;
                oif-map;
                passive;
                promiscuous-mode;
                ssm-map ssm-map-name;
```

```
                ssm-map-policy ssm-map-policy-name
                static {
                    group group {
                        source source;
                    }
                }
                version version;
            }
        }
        mld {
             interface interface-name {
                (accounting | no-accounting);
                disable;
                group-limit limit;
                group-policy;
                group-threshold value;
                immediate-leave;
                log-interval seconds;
                oif-map;
                passive;
                ssm-map ssm-map-name;
                ssm-map-policy ssm-map-policy-name;
                static {
                    group multicast-group-address {
                        exclude;
                        group-count number;
                        group-increment increment;
                        source ip-address {
                            source-count number;
                            source-increment increment;
                        }
                    }
                }
                version version;
            }
        }
        router-advertisement {
            interface interface-name {
                 current-hop-limit number;
                default-lifetime seconds;
                dns-server-address
                (managed-configuration | no-managed-configuration);
                max-advertisement-interval seconds;
```

```
                    min-advertisement-interval seconds;
                    (other-stateful-configuration | no-other-stateful-configuration);
                    prefixprefix {
                        (autonomous | no-autonomous);
                        (on-link | no-on-link);
                        preferred-lifetime seconds;
                        valid-lifetime seconds;
                    }
                    reachable-time milliseconds;
                    retransmit-timer milliseconds;
                }
            }
        }
        routing-instances routing-instance-name {
            interface interface-name;
            routing-options {
                access {
                    route prefix {
                        next-hop next-hop;
                        metric route-cost;
                        preference route-distance;
                        tag route-tag;
                        tag2 route-tag2;
                    }
                }
                access-internal {
                    route subscriber-ip-address {
                        qualified-next-hop underlying-interface {
                            mac-address address;
                        }
                    }
                }
                multicast {
                    interface interface-name {
                        no-qos-adjust;
                    }
                }
            }
            rib routing-table-name {
                access {
                    route prefix {
                        next-hop next-hop;
                        metric route-cost;
```

```
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
    }
    routing-options {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
    }
    services {
        captive-portal-content-delivery {
            auto-deactivate value;
            rule name {
                match-direction (input | input-output | output);
                term name {
```

```
                        then {
                            accept;
                            redirect url;
                            rewrite destination-address address <destination-port port-number>;
                            syslog;
                        }
                    }
                }
            }
        }
        variables {
            variable-name {
                default-value default-value;
                equals expression;
                mandatory;
                uid;
                uid-reference;
            }
        }
        version-alias profile-alias-string;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name	Name of the dynamic profile; string of up to 80 alphanumeric characters.
reassemble-packets	(Optional) Enables IPv4 reassembly of fragmented GRE packets conveyed across a soft GRE tunnel from a Wi-Fi access point to a Wi-Fi access gateway on a BNG. Reassembly is supported for fragments that range in size from 256 bytes through 8192 bytes.

NOTE:

- The maximum reassembled packet size is 13,310 bytes; this requires an MTU of 1500 bytes. The router drops reassembled packets that are larger than 13,310 bytes. The router also drops DHCP discover packets that are smaller than the MTU.
- Ordering is not maintained between fragmented packets and non-fragmented packets.
- The WAG does not support soft GRE packets with keys. Fragmented packets GRE with key are not reassembled.
- Soft GRE packet reassembly is not supported for pseudowires over redundant logical tunnels (RLT).
- The order of the last arriving fragment is not guaranteed when the reassembled packets are forwarded.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `filter, policer, hierarchical-policer, three-color-policer, and policy options` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring a Basic Dynamic Profile

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Dynamic Profiles for Subscriber Management

enable

IN THIS SECTION

- Syntax | 1863
- Hierarchy Level | 1863
- Description | 1863
- Options | 1863
- Required Privilege Level | 1864
- Release Information | 1864

Syntax

```
enable service-name {
    concurrent-data-sessions max-session-number;
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Enable the service name for the subscriber profile.

Options

service-name—Name of the enabled service.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

enable (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 1864
- Hierarchy Level | 1864
- Description | 1864
- Required Privilege Level | 1865
- Release Information | 1865

Syntax

```
enable;
```

Hierarchy Level

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

Description

Enable the Junos OS enhanced subscriber management software architecture to support configuration of dynamic interfaces and services for subscriber management. To use dynamic profiles and

authentication to create and manage dynamic subscriber interfaces and services, you must enable enhanced subscriber management and reboot the router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

> *Configuring Junos OS Enhanced Subscriber Management*
>
> *Junos OS Enhanced Subscriber Management Overview*

equals (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **1865**
- Hierarchy Level | **1866**
- Description | **1868**
- Options | **1869**
- Required Privilege Level | **1869**
- Release Information | **1869**

Syntax

```
equals {
    ascii name {
        drop drop;
```

```
        forward-only forward-only;
        local-server-group local-server-group;
    }
    hexadecimal name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
```

[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay group *name* relay-option],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay group *name* relay-option option-60],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay group *name* relay-option option-77],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay relay-option],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay group *name* relay-option],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay group *name* relay-option option-60],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay group *name* relay-option option-77],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay relay-option],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option option-60],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option option-77],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay relay-option],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option option-60],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay group *name* relay-option option-77],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay relay-option],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay relay-option option-77],

```
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Configure the exact match criteria used with the DHCP relay agent selective processing feature. DHCP relay agent compares the configured match string with the option-specific string received in DHCP client packets. If there is an exact left-to-right match, DHCP performs the action you define for the match criteria.

You can configure an unlimited number of match strings. Match strings do not support wildcard attributes.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

ascii-string	ASCII string of 1 through 255 alphanumeric characters.
hexadecimal-string	Hexadecimal string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Statement updated in Junos OS Release 17.4.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

exceed-action

IN THIS SECTION

- Syntax | 1870
- Hierarchy Level | 1870
- Description | 1870
- Required Privilege Level | 1870

Release Information | 1870

Syntax

```
exceed-action {
    drop;
    syslog;
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name max-data-sessions-
per-subscriber]
```

Description

Specify the action if the maximum data sessions per subscriber exceed the maximum limit. You must also specify the drop rate of the packets for `drop` and system log details for `syslog`.

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 11.4.

exclude (JSRC Attributes)

IN THIS SECTION

- Syntax | 1871
- Hierarchy Level | 1871
- Description | 1871
- Options | 1871
- Required Privilege Level | 1872
- Release Information | 1872

Syntax

```
exclude {
    user-name [ authorization-request | provisioning-request ];
}
```

Hierarchy Level

```
[edit access profile profile-name jsrc attributes]
```

Description

Configure the router to exclude the specified attribute-value pair (AVP) from the specified Diameter message for JSRC.

Options

- `user-name`—User-Name AVP (1). Diameter AVP name and number.
- `authorization-request`—address-authorization request in AAR message sent from JSRC to the SAE.
- `provisioning-request`—provisioning-request in AAR message sent from JSRC to the SAE.

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 14.2.

RELATED DOCUMENTATION

Excluding AVPs from Diameter Messages for JSRC | 1480

Understanding JSRC-SAE Interactions | 1469

exclude (RADIUS Attributes)

IN THIS SECTION

- Syntax | 1872
- Hierarchy Level | 1875
- Description | 1875
- Options | 1876
- Required Privilege Level | 1880
- Release Information | 1880

Syntax

```
exclude {
    acc-aggr-cir-id-asc [ access-request | accounting-start | accounting-stop ];
    acc-aggr-cir-id-bin [ access-request | accounting-start | accounting-stop ];
    acc-loop-cir-id [ access-request | accounting-start | accounting-stop ];
    acc-loop-encap [ access-request | accounting-start | accounting-stop ];
    acc-loop-remote-id [ access-request | accounting-start | accounting-stop ];
```

```
    accounting-authentic [ accounting-off | accounting-on | accounting-start | accounting-stop ]
    accounting-delay-time [ accounting-off | accounting-on | accounting-start | accounting-
stop ];
    accounting-session-id access-request;
    accounting-terminate-cause accounting-off;
    acct-request-reason [ accounting-start | accounting-stop ];
    acct-tunnel-connection [ access-request | accounting-start | accounting-stop ];
    act-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    act-data-rate-up [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    att-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    att-data-rate-up [ access-request | accounting-start | accounting-stop ];
    called-station-id [ access-request | accounting-start | accounting-stop ];
    calling-station-id [ access-request | accounting-start | accounting-stop ];
    chargeable-user-identity access-request;
    class [ accounting-start | accounting-stop ];
    cos-shaping-rate [ accounting-start | accounting-stop ];
    delegated-ipv6-prefix [ accounting-start | accounting-stop ];
    dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    dhcp-header access-request;
    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    dhcp-options [ access-request | accounting-start | accounting-stop ];
    dhcpv6-header access-request;
    dhcpv6-options [ access-request | accounting-start | accounting-stop ];
    downstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
    dsl-forum-attributes [ access-request | accounting-start | accounting-stop ];
    dsl-line-state [ access-request | accounting-start | accounting-stop ];
    dsl-type [ access-request | accounting-start | accounting-stop ];
    dynamic-iflset-name [ accounting-start | accounting-stop ];
    event-timestamp [ accounting-off | accounting-on | accounting-start | accounting-stop ];
    filter-id [ accounting-start | accounting-stop ];
    first-relay-ipv4-address [ access-request | accounting-start | accounting-stop ];
    first-relay-ipv6-address [ access-request | accounting-start | accounting-stop ];
    framed-interface-id [ access-request | accounting-start | accounting-stop ];
    framed-ip-address [ access-request | accounting-start | accounting-stop ];
    framed-ip-netmask [ access-request | accounting-start | accounting-stop ];
    framed-ip-route [ accounting-start | accounting-stop ];
    framed-ipv6-address [ access-request | accounting-start | accounting-stop ];
    framed-ipv6-pool [ accounting-start | accounting-stop ];
    framed-ipv6-prefix [ accounting-start | accounting-stop ];
    framed-ipv6-route [ accounting-start | accounting-stop ];
    framed-pool [ accounting-start | accounting-stop ]; input-ipv6-gigawords accounting-stop;
```

```
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets accounting-stop;
    input-gigawords accounting-stop;
    input-ipv6-octets accounting-stop;
    input-ipv6-packets accounting-stop;
    interface-description [ access-request | accounting-start | accounting-stop ];
    l2c-downstream-data [ access-request | accounting-start | accounting-stop ];
    l2c-upstream-data [ access-request | accounting-start | accounting-stop ];
    l2tp-rx-connect-speed [ access-request | accounting-start | accounting-stop ];
    l2tp-tx-connect-speed [ access-request | accounting-start | accounting-stop ];
    max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    max-data-rate-up [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-data-rate-up [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
    nas-identifier [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
    nas-port [ access-request | accounting-start | accounting-stop ];
    nas-port-id [ access-request | accounting-start | accounting-stop ];
    nas-port-type [ access-request | accounting-start | accounting-stop ];
    output-filter [ accounting-start | accounting-stop ];
    output-gigapackets accounting-stop;
    output-gigawords accounting-stop;
    output-ipv6-gigawords accounting-stop;
    output-ipv6-octets accounting-stop;
    output-ipv6-packets accounting-stop;
    pppoe-description [ access-request | accounting-start | accounting-stop ];
    standard-attribute number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
    }
    tunnel-assignment-id [ access-request | accounting-start | accounting-stop ];
    tunnel-client-auth-id [ access-request | accounting-start | accounting-stop ];
    tunnel-client-endpoint [ access-request | accounting-start | accounting-stop ];
    tunnel-medium-type [ access-request | accounting-start | accounting-stop ];
    tunnel-server-auth-id [ access-request | accounting-start | accounting-stop ];
    tunnel-server-endpoint [ access-request | accounting-start | accounting-stop ];
    tunnel-type [ access-request | accounting-start | accounting-stop ];
    upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
    vendor-id id-number {
```

```
        vendor-attribute vsa-number {
            packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
        }
    }
    virtual-router [ access-request | accounting-start | accounting-stop ];
}
```

Hierarchy Level

```
[edit access profile profile-name radius attributes]
```

Description

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the `ignore` statement.

You can specify attribute exclusion for multiple RADIUS message types by enclosing the message types, separated by spaces, within brackets ([ ]). You do not need brackets when specifying a single message type.

Starting in Junos OS Release 18.1R1, you can specify standard RADIUS attributes with the attribute number. You can specify VSAs with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be 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 exclude only a subset of all attributes that can be received in Access-Accept messages.

Not all attributes are available in all types of RADIUS messages.

NOTE: If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

NOTE: VSAs with dedicated option names include Juniper Networks (IANA vendor ID 4874) and DSL Forum (vendor ID 3561) VSAs.

Options

RADIUS attribute—RADIUS standard attribute or VSA:

- `acc-aggr-cir-id-asc`—Exclude Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- `acc-aggr-cir-id-bin`—Exclude Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- `acc-loop-cir-id`—Exclude Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- `acc-loop-encap`—Exclude Juniper Networks VSA 26-183, Acc-Loop-Encap.
- `acc-loop-remote-id`—Exclude Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- `accounting-authentic`—Exclude RADIUS attribute 45, Acct-Authentic.
- `accounting-delay-time`—Exclude RADIUS attribute 41, Acct-Delay-Time.
- `accounting-session-id`—Exclude RADIUS attribute 44, Acct-Session-Id.
- `accounting-terminate-cause`—Exclude RADIUS attribute 49, Acct-Terminate-Cause.
- `acct-request-reason`—Exclude Juniper Networks VSA 26-210, Acct-Request-Reason.
- `acct-tunnel-connection`—Exclude RADIUS attribute 68, Acct-Tunnel-Connection.
- `act-data-rate-dn`—Exclude Juniper Networks VSA 26-114, Act-Data-Rate-Dn.
- `act-data-rate-up`—Exclude Juniper Networks VSA 26-113, Act-Data-Rate-Up.
- `act-interlv-delay-dn`—Exclude Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn.
- `act-interlv-delay-up`—Exclude Juniper Networks VSA 26-124, Act-Interlv-Delay-Up.
- `att-data-rate-dn`—Exclude Juniper Networks VSA 26-118, Att-Data-Rate-Dn.
- `att-data-rate-up`—Exclude Juniper Networks VSA 26-117, Att-Data-Rate-Up.
- `called-station-id`—Exclude RADIUS attribute 30, Called-Station-Id.
- `calling-station-id`—Exclude RADIUS attribute 31, Calling-Station-Id.
- `chargeable-user-identity`—Exclude RADIUS attribute 89, Chargeable-User-Identity.

- `class`—Exclude RADIUS attribute 25, Class.
- `cos-shaping-rate`—Exclude Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- `delegated-ipv6-prefix`—Exclude RADIUS attribute 123, Delegated-IPv6-Prefix.
- `dhcp-gi-address`—Exclude Juniper Networks VSA 26-57, DHCP-GI-Address.
- `dhcp-header`—Exclude Juniper Networks VSA 26-208, DHCP-Header.
- `dhcp-mac-address`—Exclude Juniper Networks VSA 26-56, DHCP-MAC-Address.
- `dhcp-options`—Exclude Juniper Networks VSA 26-55, DHCP-Options.
- `dhcpv6-header`—Exclude Juniper Networks VSA 26-209, DHCPv6-Header.
- `dhcpv6-options`—Exclude Juniper Networks VSA 26-207, DHCPv6-Options.
- `dynamic-iflset-name`—Exclude Juniper Networks VSA 26-130, Qos-Set-Name.
- `downstream-calculated-qos-rate`—Exclude Juniper Networks VSA 26-141.
- `dsl-forum-attributes`—Exclude DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*.
- `dsl-line-state`—Exclude Juniper Networks VSA 26-127, DSL-Line-State.
- `dsl-type`—Exclude Juniper Networks VSA 26-128, DSL-Type.
- `event-timestamp`—Exclude RADIUS attribute 55, Event-Timestamp.
- `filter-id`—Exclude RADIUS attribute 11, Filter-Id.
- `first-relay-ipv4-address` —Exclude Juniper Networks VSA 26-189, DHCP-First-Relay-IPv4-Address.
- `first-relay-ipv6-address` —Exclude Juniper Networks VSA 26-190, DHCP-First-Relay-IPv6-Address.
- `framed-interface-id`—Exclude RADIUS attribute 96, Framed-Interface-ID.
- `framed-ip-address`—Exclude RADIUS attribute 8, Framed-IP-Address.
- `framed-ip-netmask`—Exclude RADIUS attribute 9, Framed-IP-Netmask.
- `framed-ip-route`—Exclude RADIUS attribute 22, Framed-Route.
- `framed-ipv6-address`—Exclude RADIUS attribute 168, Framed-IPv6-Address.
- `framed-ipv6-pool`—Exclude RADIUS attribute 100, Framed-IPv6-Pool.
- `framed-ipv6-prefix`—Exclude RADIUS attribute 97, Framed-IPv6-Prefix.

- framed-ipv6-route—Exclude RADIUS attribute 99, Framed-IPv6-Route.
- framed-pool—Exclude RADIUS attribute 88, Framed-Pool.
- input-filter—Exclude Juniper Networks VSA 26-10, Ingress-Policy-Name.
- input-gigapackets—Exclude Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- input-gigawords—Exclude RADIUS attribute 52, Acct-Input-Gigawords.
- input-ipv6-gigawords—Exclude Juniper Networks VSA 26-155, Acct-Input-IPv6-Gigawords.
- input-ipv6-octets—Exclude Juniper Networks VSA 26-151, Acct-Input-IPv6-Octets.
- input-ipv6-packets—Exclude Juniper Networks VSA 26-153, Acct-Input-IPv6-Packets.
- interface-description—Exclude Juniper Networks VSA 26-53, Interface-Desc.
- l2c-downstream-data—Exclude Juniper Networks VSA 26-93, L2C-Down-Stream-Data.
- l2c-upstream-data—Exclude Juniper Networks VSA 26-92, L2C-Up-Stream-Data.
- l2tp-rx-connect-speed—Exclude Juniper Networks VSA 26-163, Rx-Connect-Speed.
- l2tp-tx-connect-speed—Exclude Juniper Networks VSA 26-162, Tx-Connect-Speed.
- max-data-rate-dn—Exclude Juniper Networks VSA 26-120, Max-Data-Rate-Dn.
- max-data-rate-up—Exclude Juniper Networks VSA 26-119, Max-Data-Rate-Up.
- max-interlv-delay-dn—Exclude Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn.
- max-interlv-delay-up—Exclude Juniper Networks VSA 26-123, Max-Interlv-Delay-Up.
- min-data-rate-dn—Exclude Juniper Networks VSA 26-116, Min-Data-Rate-Dn.
- min-data-rate-up—Exclude Juniper Networks VSA 26-115, Min-Data-Rate-Up.
- min-lp-data-rate-dn—Exclude Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn.
- min-lp-data-rate-up—Exclude Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up.
- nas-identifier—Exclude RADIUS attribute 32, NAS-Identifier.
- nas-port—Exclude RADIUS attribute 5, NAS-Port.
- nas-port-id—Exclude RADIUS attribute 87, NAS-Port-Id.
- nas-port-type—Exclude RADIUS attribute 61, NAS-Port-Type.
- output-filter—Exclude Juniper Networks VSA 26-11, Egress-Policy-Name.

- `output-gigapackets`—Exclude Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- `output-gigawords`—Exclude RADIUS attribute 53, Acct-Output-Gigawords.
- `output-ipv6-gigawords`—Exclude Juniper Networks VSA 26-156, Acct-Output-IPv6-Gigawords.
- `output-ipv6-octets`—Exclude Juniper Networks VSA 26-152, Acct-Output-IPv6-Octets.
- `output-ipv6-packets`—Exclude Juniper Networks VSA 26-154, Acct-Output-IPv6-Packets.
- `packet-type`—Specify the RADIUS message type to exclude; term required when excluding a standard attribute or VSA by number rather than name. You can enclose multiple values in square brackets to specify a list of message types. Message types include Access-Request, Accounting-Off, Accounting-Off, Accounting-Start, and Accounting-Stop.
- `pppoe-description`—Exclude Juniper Networks VSA 26-24, PPPoE-Description.
- `standard-attribute` *`number`*—RADIUS standard attribute number supported by your platform. If you configure an unsupported attribute, that configuration has no effect. When you use this option, you must use the `packet-type` term to specify the message from which the attribute is excluded.
- `tunnel-assignment-id`—Exclude RADIUS attribute 82, Tunnel-Assignment-ID.
- `tunnel-client-auth-id`—Exclude RADIUS attribute 90. Tunnel-Client-Auth-ID.
- `tunnel-client-endpoint`—Exclude RADIUS attribute 66, Tunnel-Client-Endpoint.
- `tunnel-medium-type`—Exclude RADIUS attribute 65, Tunnel-Medium-Type.
- `tunnel-server-auth-id`—Exclude RADIUS attribute 91, Tunnel-Server-Auth-ID.
- `tunnel-server-endpoint`—Exclude RADIUS attribute 67, Tunnel-Server-Endpoint.
- `tunnel-type`—Exclude RADIUS attribute 64, Tunnel-Type.
- `upstream-calculated-qos-rate`—Exclude Juniper Networks VSA 26-142
- `vendor-attribute` *`vsa-number`*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. If you configure an unsupported VSA, that configuration has no effect. When you use this option, you must use the `packet-type` term to specify the message from which the attribute is excluded.
- `vendor-id` *`id-number`*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect.
- `virtual-router`—Exclude Juniper Networks VSA 26-1.

RADIUS message type:

- access-request—RADIUS Access-Request messages.
- accounting-off—RADIUS Accounting-Off messages.
- accounting-on—RADIUS Accounting-On messages.
- accounting-start—RADIUS Accounting-Start messages.
- accounting-stop—RADIUS Accounting-Stop messages.

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 9.1.

downstream-calculated-qos-rate, dsl-forum-attributes, and upstream-calculated-qos-rate options added in Junos OS Release 11.4.

cos-shaping-rate and filter-id options added in Junos OS Release 13.2.

pppoe-description option added in Junos OS Release 14.2.

virtual-router option added in Junos OS Release 15.1.

first-relay-ipv4-address and first-relay-ipv6-address options added in Junos OS Release 16.1.

acc-loop-encap and acc-loop-remote-id options added in Junos OS Release 16.1R4.

access-request option support for all tunnel attributes added in Junos OS Release 15.1R7, 16.1R5, 16.2R2, 17.1R2, 17.2R2, and 17.3R1 for MX Series.

packet-type, standard-attribute, vendor-attribute, and vendor-id options added in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

Standard and Vendor-Specific RADIUS Attributes | 3

excluded-address (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1881
- Hierarchy Level | 1881
- Description | 1881
- Options | 1881
- Required Privilege Level | 1882
- Release Information | 1882

Syntax

```
excluded-address ip-address;
```

Hierarchy Level

```
[edit access address-assignment pool name family (inet | inet6)],
[edit logical-systems name access address-assignment pool name family (inet | inet6)],
[edit logical-systems name routing-instances name access address-assignment pool name family (inet
| inet6)],
[edit routing-instances name access address-assignment pool name family (inet | inet6)]
```

Description

Specify an address to exclude from consideration when addresses are allocated from the corresponding address pool. If an address that you configure for exclusion has already been allocated, the subscriber that has that address is logged out. The address is then deallocated and marked for exclusion from future allocation.

Options

ip-address IPv4 or IPv6 address to exclude from the address pool for the specified family.

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 18.1R1.

RELATED DOCUMENTATION

- Preventing Addresses from Being Allocated from an Address Pool | 1149
- Address-Assignment Pools Overview | 1139
- Address-Assignment Pool Configuration Overview | 1146

excluded-range (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1882
- Hierarchy Level | 1883
- Description | 1883
- Options | 1883
- Required Privilege Level | 1883
- Release Information | 1883

Syntax

```
excluded-range name low minimum-value high maximum-value;
```

Hierarchy Level

```
[edit access address-assignment pool name family (inet | inet6)],
[edit logical-systems name access address-assignment pool name family (inet | inet6)],
[edit logical-systems name routing-instances name access address-assignment pool name family (inet
| inet6)],
[edit routing-instances name access address-assignment pool name family (inet | inet6)]
```

Description

Specify a range of consecutive addresses to exclude from consideration when addresses are allocated from the corresponding address pool. For example, you might want to reserve certain addresses or ranges to be used only for static subscribers. If a range you configure for exclusion includes an address that has already been allocated, the subscriber that has that address is logged out. The address is then deallocated and marked for exclusion from future allocation.

Options

high Upper limit of excluded range of addresses.

low Lower limit of range of addresses.

name Name of a range of addresses to exclude from the address pool for the specified family.

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 18.1R1.

RELATED DOCUMENTATION

Preventing Addresses from Being Allocated from an Address Pool | 1149

Address-Assignment Pools Overview | 1139

Address-Assignment Pool Configuration Overview | 1146

external-authority

IN THIS SECTION

- Syntax | 1884
- Hierarchy Level | 1884
- Description | 1884
- Required Privilege Level | 1884
- Release Information | 1885

Syntax

```
external-authority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Specify that an external authority (for example, RADIUS or Diameter) provides the address assignment.

When RADIUS is the external authority, the router uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. When Diameter is the external authority, the router uses the Diameter counterpart of RADIUS Framed-IPv6-Pool attribute.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **773**
- Understanding Differences Between Legacy DHCP and Extended DHCP | **691**
- Address-Assignment Pools Overview | **1139**

failover (System Process)

IN THIS SECTION

- Syntax | **1885**
- Hierarchy Level | **1885**
- Description | **1886**
- Options | **1886**
- Required Privilege Level | **1886**
- Release Information | **1886**

Syntax

```
failover (alternate-media | other-routing-engine);
```

Hierarchy Level

```
[edit system processes process-name]
```

Description

Configure the router to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.

Options

process-name—Junos OS process name. Some of the processes that support the `failover` statement are `bootp`, `chassis-control`, `craft-control`, `ethernet-connectivity-fault-management`, `init`, `interface-control`, `neighbor-liveness`, `pfe`, `redundancy-interface-process`, `routing`, `smg-service`, and `vrrp`.

`alternate-media`—Use the Junos OS image on alternate media during the reboot.

`other-routing-engine`—On routers with dual Routing Engines, use the Junos OS image on the other Routing Engine during the reboot. That Routing Engine assumes the primary role; in the usual configuration, the other Routing Engine is the designated backup Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

When a Software Process Fails

processes

family (Address-Assignment Pools)

IN THIS SECTION

Syntax | 1887

- Hierarchy Level | **1887**
- Description | **1887**
- Options | **1888**
- Required Privilege Level | **1888**
- Release Information | **1888**

Syntax

```
family family {
    dhcp-attributes {
        [protocol-specific attributes]
    }
    excluded-address ip-address;
    excluded-range name low minimum-value high maximum-value;
    host hostname {
        hardware-address mac-address;
        ip-address ip-address;
    }
    network ip-prefix/<prefix-length>;
    prefix ipv6-prefix;
    range range-name {
        high upper-limit;
        low lower-limit;
        prefix-length prefix-length;
    }
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name]
```

Description

Configure the protocol family for the address-assignment pool.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

family—Protocol family:

- `inet`—Internet Protocol version 4 suite
- `inet6`—Internet Protocol version 6 suite

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools Overview | 1139

Address-Assignment Pool Configuration Overview | 1146

family-state-change-immediate-update

IN THIS SECTION

Syntax | 1889

Hierarchy Level | 1889

- Description | 1889
- Required Privilege Level | 1889
- Release Information | 1889

Syntax

```
family-state-change-immediate-update family-state-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Send an Acct-Update message to notify address family activation state change.

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 17.4.

RELATED DOCUMENTATION

Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview | 843

final-response-timeout (OCS Partition)

IN THIS SECTION

- Syntax | 1890
- Hierarchy Level | 1890
- Description | 1890
- Options | 1891
- Required Privilege Level | 1891
- Release Information | 1891

Syntax

```
final-response-timeout seconds;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the amount of time in seconds before an Online Charging System (OCS) partition stops attempting to send the final interrogation during the subscriber logout process. When a subscriber starts to logs out, the OCS sends a final interrogation. If there is no response within 24 hours, then the system continues and starts sending a logout message to the Policy and Charging Rule Function (PCRF). If there is no response within the next 24 hours, then the subscriber logout proceeds.

NOTE: Any configuration changes made to this statement apply to all subscribers currently waiting to log out within a 60 second period.

Options

seconds Number of seconds to wait before an OCS partition stops sending the final interrogation using a CCR-GY-T message.

- **Default:** 7200
- **Range:** 0 through 86,400 seconds (24 hours)

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 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition | **1449**

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gy Interactions Between the Router and the OCS | **1431**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

force-continue (OCS Partition)

IN THIS SECTION

- Syntax | **1892**
- Hierarchy Level | **1892**
- Description | **1892**
- Required Privilege Level | **1892**

Release Information | 1892

Syntax

```
force-continue;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Enable subscriber traffic to flow before the first interrogation with the Online Charging System (OCS) occurs.

Wireline customers often control user services solely through the Policy Control and Charging Rules Function (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 to force the broadband Policy and Charging Enforcement Function (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.

NOTE: The `force-continue` state is required; you must configure it in the OCS partition.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

forward-only (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 1893
- Hierarchy Level | 1893
- Description | 1894
- Required Privilege Level | 1894
- Release Information | 1894

Syntax

```
forward-only;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with)],
[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 ...]
```

Description

Forward specified DHCP client packets, without creating a new subscriber session, when you use DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent. Use this command to manage subscribers, when large number of DHCP relay requests are received and license limitations are exceeded. This command uses option 82 and applies "scale-subscriber" license.

NOTE: DHCP packets forwarded with the `forward-only` statement do not consider other configurations except for the `trust-option-82` option. The DHCP relay agent ignores all other configured options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

forward-only (DHCP Relay Agent)

IN THIS SECTION

Syntax | 1895

- Hierarchy Level | **1895**
- Description | **1895**
- Default | **1895**
- Options | **1896**
- Required Privilege Level | **1896**
- Release Information | **1896**

Syntax

```
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name 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 ...]
```

Description

Specify the VRF location of the DHCP server when configuring secure DHCP traffic between the DHCP server and DHCP client when the two reside in different VRFs.

Default

Logical system and routing instance from where the configuration is applied.

Options

logical-system (Optional) Logical system in which the DHCP server resides.

- `current`—Logical system from which the configuration is applied.
- `default`—Root logical system.
- *`logical-system-name`*—A specific logical system.

routing-instance (Optional) Routing instance in which the DHCP server resides.

- `current`—Routing instance from which the configuration is applied.
- `default`—Root routing instance.
- *`logical-system-name`*—A specific routing instance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | **739**

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | **740**

forward-only-replies (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1897
- Hierarchy Level | 1897
- Description | 1897
- Required Privilege Level | 1897
- Release Information | 1898

Syntax

```
forward-only-replies;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that reply packets are for the forward-only support that is configured in option 82 interface ID of the reply packet. You must configure this statement for forward-only support when the client and server are in different logical system/routing instances.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | 739

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | 740

forwarding (Diameter Network Element)

IN THIS SECTION

- Syntax | 1898
- Hierarchy Level | 1899
- Description | 1899
- Required Privilege Level | 1899
- Release Information | 1899

Syntax

```
forwarding {
    route dne-route-name {
        destination realm realm-name <host hostname>;
        function function-name <partition partition-name>;
        metric route-metric;
    }
}
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Define the criteria that specify which destinations are reachable through the Diameter network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

function (Diameter Network Element)

IN THIS SECTION

- Syntax | 1900
- Hierarchy Level | 1900
- Description | 1900
- Default | 1900

Options | 1900
Required Privilege Level | 1901
Release Information | 1901

Syntax

```
function [function-name];
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Specify one or more applications (function) associated with a Diameter network element.

Default

By default, all functions are associated with (supported by) the network element.

Options

function-name—Application (function) associated with the route. You can list multiple functions with the element.

- `gx-plus`—Associate the Gx-Plus application with the network element.
- `jsrc`—Associate the JSRC application with the network element.
- `nasreq`—Associate the NASREQ application with the network element.
- `packet-triggered-subscribers`—Associate the packet-triggered subscribers application with the network element.

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 10.1.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

Support for NASREQ added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

function (Diameter Route)

IN THIS SECTION

- Syntax | 1902
- Hierarchy Level | 1902
- Description | 1902
- Default | 1902
- Options | 1902
- Required Privilege Level | 1902
- Release Information | 1902

Syntax

```
function function-name <partition partition-name>;
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding route dne-route-name]
```

Description

Specify the application (function) associated with a destination and metric. Together, these three elements define a route reachable through a Diameter network element.

Default

All functions are associated with the route.

Options

function-name—Application (function) associated with the route. Gx-Plus, JSRC, NASREQ, and packet-triggered-subscribers are the applications currently supported.

`partition` *partition-name*—(Optional) Partition associated with the application (function).

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 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

Support for NASREQ added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

ggsn-address (OCS Partition)

IN THIS SECTION

- Syntax | 1903
- Hierarchy Level | 1903
- Description | 1903
- Options | 1903
- Required Privilege Level | 1904
- Release Information | 1904

Syntax

```
ggsn-address address;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the GGSN-Address AVP value in all CCR-GY messages.

Options

address Address of the GGSN-Address AVP value to include in all CCR-GY messages.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

ggsn-mcc-mnc (OCS Partition)

IN THIS SECTION

- Syntax | 1904
- Hierarchy Level | 1905
- Description | 1905
- Options | 1905
- Required Privilege Level | 1905
- Release Information | 1905

Syntax

```
ggsn-mcc-mnc ggsn-mcc-mnc;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the value of the 3GPP-GGSN-MCC-MNC AVP value to include in all CCR-GY messages. The value is the mobile country code (MCC) and mobile network code (MNC) of the network that the GGSN belongs to.The combined MCC and MNC uniquely identify the mobile network operator.

Options

ggsn-mcc-mnc Value of the mobile network country and network codes to include in all CCR-GY messages.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

global (Gx-Plus)

IN THIS SECTION

- Syntax | 1906
- Hierarchy Level | 1906
- Description | 1906
- Required Privilege Level | 1906
- Release Information | 1907

Syntax

```
global {
    include-ipv6;
    max-outstanding-requests number;
}
```

Hierarchy Level

```
[edit access gx-plus]
```

Description

Configure global attributes for the Gx-Plus application.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus | 1404

global (OCS)

IN THIS SECTION

- Syntax | 1907
- Hierarchy Level | 1907
- Description | 1907
- Required Privilege Level | 1908
- Release Information | 1908

Syntax

```
global {
    service-context-id service-context ;
}
```

Hierarchy Level

```
[edit access ocs]
```

Description

Configure global attributes and data elements 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). 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.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 16.2.

RELATED DOCUMENTATION

- Configuring OCS Global Parameters | 1462
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

global (PCRF)

IN THIS SECTION

- Syntax | 1909
- Hierarchy Level | 1909
- Description | 1909
- Required Privilege Level | 1909
- Release Information | 1909

Syntax

```
global {
    rule-param avp-code;
}
```

Hierarchy Level

```
[edit access pcrf]
```

Description

Configure global attributes and data elements 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). 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.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

group (DHCP Local Server)

IN THIS SECTION

- Syntax | 1910
- Hierarchy Level | 1914
- Description | 1914
- Options | 1914
- Required Privilege Level | 1914
- Release Information | 1915

Syntax

```
group group-name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            relay-agent-interface-id
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-profile-
name>;
```

```
interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82>;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
```

```
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            rapid-commit;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        trace;
        upto upto-interface-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82>;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
```

```
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    reconfigure {
```

```
        attempts attempt-count;
        clear-on-terminate;
        strict;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Options

group-name—Name of the group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Understanding Differences Between Legacy DHCP and Extended DHCP | 691
- Grouping Interfaces with Common DHCP Configurations | 852
- Specifying Authentication Support | 833
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

group (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1915
- Hierarchy Level | 1920
- Description | 1920
- Options | 1920
- Required Privilege Level | 1920
- Release Information | 1920

Syntax

```
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
```

```
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    interface interface-name {
        access-profile profile-name;
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
```

```
                }
            }
        }
        overrides {
            allow-no-end-option;
            allow-snooped-clients;
            always-write-giaddr;
            always-write-option-82;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match <option60-and-option82 | incoming-interface>;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            disable-relay;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            layer2-unicast-replies;
            no-allow-snooped-clients;
            no-bind-on-request;
            proxy-mode;
            relay-source
            replace-ip-source-with;
            send-release-on-delete;
            trust-option-82;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        trace;
        upto upto-interface-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
```

```
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82>;
        client-negotiation-match incoming-interface;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
```

```
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
        }
        remote-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        server-id-override
    }
    remote-id-mismatch disconnect;
    route-suppression;
```

```
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Extended DHCP Relay Agent Overview | 695
- Configuring DHCP Relay Agent
- Configuring Group-Specific DHCP Relay Options | 857
- Grouping Interfaces with Common DHCP Configurations | 852
- Specifying Authentication Support | 833
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

group (Static Subscribers)

IN THIS SECTION

- Syntax | 1921
- Hierarchy Level | 1922
- Description | 1922
- Options | 1922
- Required Privilege Level | 1922
- Release Information | 1923

Syntax

```
group group-name {
    access-profile profile-name;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
    }
    authentication {
        password password-string;
        username-include {
            delimiter delimiter-character;
            domain-name domain-name;
            interface;
```

```
            logical-system-name;
            routing-instance-name;
             user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    interface  interface-name  <exclude> <upto upto-interface-name>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit system services static-subscribers]
```

Description

Configure a static subscriber group with values that override the values configured at the `[edit system services static-subscribers]` hierarchy level for subscribers outside the group. Includes the subscriber access and dynamic profiles, the authentication parameters that trigger the Access-Request message to AAA for static subscribers in the group, and the statically configured interfaces that form the group.

NOTE: The logical system and routing instance in which the group is configured must match the logical system and routing instance where the static interfaces are configured.

Options

group-name—Name of a group that defines authentication parameters for static subscribers to override the global authentication configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview | 1487

Creating a Static Subscriber Group | 1494

gsmp-syn-timeout (ANCP)

IN THIS SECTION

- Syntax | 1923
- Hierarchy Level | 1923
- Description | 1924
- Options | 1924
- Required Privilege Level | 1924
- Release Information | 1924

Syntax

```
gsmp-syn-timeout seconds;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure the maximum period that the ANCP agent waits before sending a SYN message to an ANCP neighbor to negotiate the adjacency. If the neighbor sends a SYN message during this period, the ANCP agent uses the partition information in the neighbor's message when generating its own initial SYN message to the neighbor. The agent does not wait for the period to expire if it receives a SYN message from the neighbor.

Options

seconds—Number of seconds the ANCP agent waits.

- **Range:** 1 through 60 seconds
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring the ANCP Agent | **1253**

Configuring the ANCP Agent to Learn ANCP Partition IDs | **1260**

gsmp-syn-wait (ANCP)

IN THIS SECTION

- Syntax | **1925**
- Hierarchy Level | **1925**

- Description | 1925
- Default | 1925
- Required Privilege Level | 1925
- Release Information | 1926

Syntax

```
gsmp-syn-wait;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Enable the ANCP agent to learn partition ID information from neighbors, in support of nonzero ANCP partition IDs. This statement forces the ANCP agent to delay sending a SYN message during adjacency negotiation for a configurable period. When the neighbor sends a SYN message to the ANCP agent during that period, the agent learns the partition ID information from the neighbor and uses that information when it sends its own SYN message. If the agent does not receive the message during the period, then it sends a SYN message to the neighbor when the period times out.

Default

This statement is disabled. The ANCP agent does not wait before sending the initial SYN message and does not support nonzero partition IDs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3

RELATED DOCUMENTATION

- Configuring the ANCP Agent | 1253
- Configuring the ANCP Agent to Learn ANCP Partition IDs | 1260

gx-plus (Gx-Plus)

IN THIS SECTION

- Syntax | 1926
- Hierarchy Level | 1927
- Description | 1927
- Required Privilege Level | 1927
- Release Information | 1927

Syntax

```
gx-plus {
    global {
        include-ipv6;
        max-outstanding-requests number;
    }
    partition partition-name  {
        diameter-instance instance-name;
        destination-host hostname;
        destination-realm realm;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Gx-Plus application to interact with a PCRF to authorize and provision subscribers.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus | 1404

host (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1928
- Hierarchy Level | 1928
- Description | 1928
- Options | 1928
- Required Privilege Level | 1928

- Release Information | 1929

Syntax

```
host hostname {
    hardware-address mac-address;
    ip-address ip-address;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Description

Configure a static binding for the specified client.

Options

hostname	Name of the client.
hardware-address *mac-address*	Specify the MAC address of the client. This is the hardware address that identifies the client on the network. • *mac-address*—MAC address of the client.
ip-address *ip-address*	Specify the reserved IP address assigned to the client. • *ip-address*—IP version 4 (IPv4) address.

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 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

host-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1929
- Hierarchy Level | 1929
- Description | 1929
- Required Privilege Level | 1930
- Release Information | 1930

Syntax

```
host-name name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82]
```

Description

Supports the addition of vendor-specific hostname in the option 82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The hostname can be a string of characters such as **Juniper-AB-1**.

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The hostname is option-data 1 (the location is option-data 2). The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

host-name (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | **1930**
- Hierarchy Level | **1931**
- Description | **1931**
- Required Privilege Level | **1931**
- Release Information | **1931**

Syntax

```
host-name name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-option-vendor-specific],
[edit forwarding-options dhcp-relaygroup group-name dhcpv6 relay-option-vendor-specific]
```

Description

Supports the addition of vendor-specific hostname in the vendor-specific option (17) of DHCPv6 control messages on server-facing interfaces. The hostname can be a string of characters such as **Juniper-AB-1**.

Junos automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The hostname is option-data 1 (the location is option-data 2). The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

ietf-mode

IN THIS SECTION

- Syntax | **1932**
- Hierarchy Level | **1932**

- Description | 1932
- Default | 1932
- Required Privilege Level | 1932
- Release Information | 1932

Syntax

```
ietf-mode
```

Hierarchy Level

```
[edit protocols ancp neighbor ip-address]
```

Description

Configure the ANCP agent to run in a mode that is not backward compatible with Internet draft-wadhwa-gsmp-l2control-configuration-00.txt, *GSMP extensions for layer2 control (L2C)*. Include this statement when pre-ietf mode has been configured globally for the ANCP agent, but you want one or more neighbors to run in the default mode.

Default

ANCP does not run in a backward-compatible mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Configuring ANCP Neighbors | 1255

immediate-update

IN THIS SECTION

- Syntax | 1933
- Hierarchy Level | 1933
- Description | 1933
- Required Privilege Level | 1933
- Release Information | 1934

Syntax

```
immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router or switch to send an Acct-Update message to the RADIUS accounting server on receipt of a response (for example, an ACK or timeout) to the Acct-Start message.

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 9.1.

RELATED DOCUMENTATION

- RADIUS Servers and Parameters for Subscriber Access | 476
- RADIUS Accounting for Subscriber Access | 570

include-ipv6 (Gx-Plus)

IN THIS SECTION

- Syntax | 1934
- Hierarchy Level | 1934
- Description | 1934
- Default | 1935
- Required Privilege Level | 1935
- Release Information | 1935

Syntax

```
include-ipv6;
```

Hierarchy Level

```
[edit access gx-plus global]
```

Description

Include IPv6 subscribers in Gx-Plus provisioning requests.

Default

By default, IPv6 subscribers are not included.

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 12.3.

RELATED DOCUMENTATION

Configuring Gx-Plus Global Attributes | 1406

Configuring Gx-Plus | 1404

include-irb-and-l2

IN THIS SECTION

- Syntax | 1935
- Hierarchy Level | 1936
- Description | 1936
- Required Privilege Level | 1938
- Release Information | 1938

Syntax

```
include-irb-and-l2;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Include both the integrated routing and bridging (IRB) interface name and Layer 2 interface name in the `circuit-id` or `remote-id` value in the DHCP option 82 information. VLAN tags are global.

For leasequery and bulk leasequery operations that involve 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.

When you configure the `include-irb-and-l2` statement without including the `no-vlan-interface` statement, the format is as follows:

- Bridge domain:

  ```
  (fe | ge)-fpc/pic/port.subunit:bridge-domain-name+irb.subunit
  ```

- VLAN:

  ```
  (fe | ge)-fpc/pic/port.subunit:vlan-name+irb.subunit
  ```

NOTE: For remote systems, the *subunit* is required and is used to differentiate an interface.

When you configure both the `include-irb-and-l2` statement and the `use-vlan-id` statement, the format is as follows:

```
(fe | ge)-fpc/pic/port.subunit:svlan-id-vlan-id+irb.subunit
```

NOTE: The *svlan-id-vlan-id* represents the VLANs associated with the bridge domain.

When you configure both the `include-irb-and-l2` and `no-vlan-interface-name` statements, the format is as follows:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

When you configure both the `include-irb-and-l2` and `use-interface-description` statements, the format displays the description for the Layer 2 interface:

```
l2_descr:vlan-name+irb.subunit
```

If you configure both the `include-irb-and-l2` and `use-interface-description` statements, and no description for the Layer 2 interface is found, the format displays the Layer 2 logical interface name:

```
(fe | ge)-fpc/pic/port.subunit:vlan-name+irb.subunit
```

When you configure the `include-irb-and-l2` statement with both the `no-vlan-interface-name` and `use-interface-description` statements, the format displays as follows:

```
l2_descr+irb.subunit
```

If you configure the include-irb-and-l2 statement with both the no-vlan-interface-name and use-interface-description statements, and no description is found for the Layer 2 interface, the format displays as follows:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

NOTE: The EX Series switches that support the include-irb-and-l2 statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

- Using DHCP Relay Agent Option 82 Information | 751
- Using DHCP Relay Agent Option 82 Information | 751
- Configuring DHCPv6 Relay Agent Options | 917

include-option-82 (DHCP Local Server)

IN THIS SECTION

- Syntax | 1939
- Hierarchy Level | 1939
- Description | 1939
- Options | 1939

Required Privilege Level | 1940
Release Information | 1940

Syntax

```
include-option-82 {
    forcerenew;
    nak;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server (dhcpv6) group group-name overrides],
[edit system services dhcp-local-server (dhcpv6) group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server (dhcpv6) ...overrides],
[edit logical-systems logical-system-name system services dhcp-local-server
(dhcpv6) ...overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server
(dhcpv6) ...overrides]
```

Description

Specify that the DHCP server include option 82 information in NAK and forcerenew messages when you configure secure communications between the DHCP server and DHCP clients that are in different VRFs. You can configure support globally, for a group of interfaces, or for a specific interface.

Options

forcerenew	Include option 82 in DHCP forcerenew messages.
nak	Include option 82 in DHCP NAK messages.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs | 739

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances | 740

inet (Interfaces)

IN THIS SECTION

- Syntax | 1940
- Hierarchy Level | 1941
- Description | 1941
- Options | 1941
- Required Privilege Level | 1942
- Release Information | 1942

Syntax

```
inet {
    address source;
    auto-configure {
        address-ranges {
```

```
            authentication {
                password password-string;
                username-include {
                    auth-server-realm realm-string;
                    delimiter delimiter-character;
                    domain-name domain-name;
                    interface-name;
                    source-address;
                    user-prefix user-prefix-string;
                }
            }
            dynamic-profile profile-name {
                network ip-address {
                    range name {
                        low lower-limit;
                        high upper-limit;
                    }
                }
            }
        }

    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux]
```

Description

Specify the inet family for the demultiplexing (demux) interface options.

Options

address-source *address* Specify the source IPv4 or IPv6 address or prefix value from which to inherit configuration data for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

inet6 (Interfaces)

IN THIS SECTION

- Syntax | 1942
- Hierarchy Level | 1943
- Description | 1943
- Options | 1943
- Required Privilege Level | 1943
- Release Information | 1944

Syntax

```
inet6 {
    address source;
    auto-configure {
        address-ranges {
            authentication {
                password password-string;
```

```
                username-include {
                    auth-server-realm realm-string;
                    delimiter delimiter-character;
                    domain-name domain-name;
                    interface-name;
                    source-address;
                    user-prefix user-prefix-string;
                }
            }
            dynamic-profile profile-name {
                network ip-address {
                    range name {
                        low lower-limit;
                        high upper-limit;
                    }
                }
            }
        }

    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux]
```

Description

Specify the inet6 family for the demultiplexing (demux) interface options.

Options

address-source *address* Specify the source IPv4 or IPv6 address or prefix value from which to inherit configuration data for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

interface (DHCP Local Server)

IN THIS SECTION

- Syntax | 1944
- Hierarchy Level | 1946
- Description | 1946
- Options | 1946
- Required Privilege Level | 1946
- Release Information | 1947

Syntax

```
interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82 | incoming-interface>;
        client-negotiation-match  incoming-interface;
        delay-advertise {
```

```
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server 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. Also, you cannot use an interface that is being used by the DHCP relay agent.

NOTE: DHCP values are supported in integrated routing and bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently supports only static DHCP configurations.

Options

`exclude`—Exclude an interface or a range of interfaces from the group. This option and the `overrides` option are mutually exclusive.

interface-name—Name of the interface. You can repeat this option multiple times.

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the *upto-interface-name* must be the same as the device name of the *interface-name*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Options `upto` and `exclude` introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

- Understanding Differences Between Legacy DHCP and Extended DHCP | 691
- Grouping Interfaces with Common DHCP Configurations | 852
- Specifying Authentication Support | 833

interface (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1947
- Hierarchy Level | 1948
- Description | 1948
- Options | 1949
- Required Privilege Level | 1949
- Release Information | 1949

Syntax

```
interface dhcp-interface-name {
    access-profile profile-name;
    exclude;
    overrides {
        allow-no-end-option
        allow-snooped-clients;
        always-write-giaddr;
```

```
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82 | incoming-interface>;
        client-negotiation-match incoming-interface;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent 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. Also, you cannot use an interface that is being used by the DHCP local server. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

NOTE: DHCP values are supported in integrated routing and bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. .

Options

`exclude`—Exclude an interface or a range of interfaces from the group. This option and the `overrides` option are mutually exclusive.

`interface-name`—Name of the interface. You can repeat this option multiple times.

`overrides`—Override the specified default configuration settings for the interface. The `overrides` statement is described separately.

`upto-interface-name`—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the *`upto-interface-name`* must be the same as the device name of the *`interface-name`*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Options `upto` and `exclude` introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Grouping Interfaces with Common DHCP Configurations | 852

Specifying Authentication Support | 833

interface (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 1950
- Hierarchy Level | 1951
- Description | 1951
- Options | 1951
- Required Privilege Level | 1951
- Release Information | 1951

Syntax

```
interface interface-name {
     current-hop-limitnumber;
    default-lifetime seconds;
    dns-server-address
    (managed-configuration | no-managed-configuration);
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
    (other-stateful-configuration | no-other-stateful-configuration);
    prefixprefix {
        (autonomous | no-autonomous);
        (on-link | no-on-link);
        preferred-lifetime seconds;
        valid-lifetime seconds;
    }
    reachable-time milliseconds;
    retransmit-timer milliseconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement]
```

Description

Dynamically configure router advertisement properties on an interface. To dynamically configure interface properties, include the *$junos-interface-name* dynamic variable for the interface name.

Options

interface-name—Name of an interface. Specify the *$junos-interface-name* dynamic variable or the full, static interface name, including the physical and logical address components.

NOTE: Even though you can specify the static interface name when defining the interface, we recommend using dynamic variable when configuring this statement.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

interface (Static Subscriber Group)

IN THIS SECTION

- Syntax | 1952
- Hierarchy Level | 1952
- Description | 1952
- Options | 1953
- Required Privilege Level | 1953
- Release Information | 1953

Syntax

```
interface interface-name <exclude> <upto upto-interface-name>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers group group-name]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which static subscribers are created. You can repeat the `interface interface-name` statement to specify multiple interfaces within a group. You must configure each interface in only one group.

NOTE: The logical system and routing instance in which the static interfaces are configured must match the logical system and routing instance where the group is configured.

Options

`exclude`—(Optional) Exclude an interface or a range of interfaces from the group.

`interface-name`—Name of the interface on which static subscribers are created. If you do not specify a unit number for the interface, then .0 is assumed. For example, `ge-0/1/0` is interpreted as `ge-0/1/0.0`.

`upto-interface-name`—(Optional) The upper end of the range of interfaces; the lower end of the range is the *interface-name* entry. The interface device name of *upto-interface-name* must be the same as the device name of *interface-name*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for IPv6 and IPv4 demux static interfaces introduced in Junos OS Release 11.2.

Support for pseudowire interfaces over logical tunnels introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview | 1487

Creating a Static Subscriber Group | 1494

interface (Static Subscriber Username)

IN THIS SECTION

- Syntax | 1954
- Hierarchy Level | 1954
- Description | 1955
- Required Privilege Level | 1955
- Release Information | 1955

Syntax

```
interface;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers authentication
username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that a modified version of the interface name is included as part of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message. The interface name is modified by replacing the "/" character with the "-" character. For example, ge-0/1/2.50 is converted to ge-0-1-2.50.

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487
- Configuring the Static Subscriber Global Username | 1492
- Configuring the Static Subscriber Group Username | 1497

interface-client-limit (DHCP Local Server)

IN THIS SECTION

- Syntax | 1956
- Hierarchy Level | 1956
- Description | 1957
- Default | 1957
- Options | 1957
- Required Privilege Level | 1957
- Release Information | 1957

Syntax

```
interface-client-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
overrides],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group interface interface-name group-name
overrides],
```

```
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
```

Description

Set the maximum number of DHCP subscribers or DHCP clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group.

NOTE: Do not use this statement for dual-stack subscribers. Instead, use the "dual-stack-interface-client-limit" on page 1833 statement for dual-stack subscribers.

Default

No limit

Options

number—Maximum number of clients allowed.

- **Range:** 1 through 500,000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Specifying the Maximum Number of DHCP Clients Per Interface | 862

Overriding the Default DHCP Local Server Configuration Settings | 707

interface-client-limit (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1958
- Hierarchy Level | 1958
- Description | 1959
- Default | 1959
- Options | 1959
- Required Privilege Level | 1959
- Release Information | 1960

Syntax

```
interface-client-limit number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Set the maximum number of DHCP (or DHCPv6) subscribers or clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

NOTE: Do not use this statement for dual-stack subscribers. Instead, use the "dual-stack-interface-client-limit" on page 1833 statement for dual-stack subscribers.

Default

No limit

Options

number—Maximum number of clients allowed.

- **Range:** 1 through 500,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- dhcp-relay | 1752
- Extended DHCP Relay Agent Overview | 695
- Configuring Group-Specific DHCP Relay Options | 857
- Overriding the Default DHCP Relay Configuration Settings | 709

interface-delete (Subscriber Management or DHCP Client Management)

IN THIS SECTION

- Syntax | 1960
- Hierarchy Level | 1960
- Description | 1961
- Required Privilege Level | 1961
- Release Information | 1961

Syntax

```
interface-delete;
```

Hierarchy Level

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

Description

On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.

On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | **866**

interface-description (DHCP Local Server)

IN THIS SECTION

- Syntax | **1962**
- Hierarchy Level | **1962**
- Description | **1962**
- Options | **1962**
- Required Privilege Level | **1963**
- Release Information | **1963**

Syntax

```
interface-description (device-interface | logical-interface);
```

Hierarchy Level

```
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify that the interface description for the device (physical) interface or the logical interface is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

The interface description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. The delimiter is configured with the `delimiter` statement at the `[edit system services dhcp-local-server ... username-include]` hierarchy level.

NOTE: The username delimiter must not be a character that is part of the interface description. For example, if the description includes a forward slash (/), then the delimiter must not be a forward slash. The default delimiter is a period.

Options

device-interface Use the text description configured for the device (physical) interface.

logical-interface Use the text description configured for the logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- Creating Unique Usernames for DHCP Clients | 834

interface-description (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1963
- Hierarchy Level | 1964
- Description | 1964
- Options | 1964
- Required Privilege Level | 1964
- Release Information | 1965

Syntax

```
interface-description (device-interface | logical-interface);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[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 ...]
```

Description

Specify that the interface description for the device (physical) interface or the logical interface is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

The interface description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. The delimiter is configured with the `delimiter` statement at the `[edit forwarding-options dhcp-relay ... username-include]` hierarchy level.

NOTE: The username delimiter must not be a character that is part of the interface description. For example, if the description includes a forward slash (/), then the delimiter must not be a forward slash. The default delimiter is a period.

Options

device-interface Description of the physical interface.

logical-interface Description of the logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- Creating Unique Usernames for DHCP Clients | 834

interface-description-format

IN THIS SECTION

- Syntax | 1965
- Hierarchy Level | 1965
- Description | 1966
- Options | 1966
- Required Privilege Level | 1966
- Release Information | 1967

Syntax

```
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the information that is excluded from the interface description that the device passes to RADIUS for inclusion in the RADIUS attributes such as NAS-Port-ID (87) or Calling-Station-ID (31).

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. If you exclude the subinterface, the description becomes ge-1/2/0:100.

The default format for channelized interfaces is as follows:

interface-type-slot/adapter/channel.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. If you exclude the channel information, the description becomes xe-0/1/2.4-5.6.

Options

exclude-adapter —(Optional) Exclude the adapter from the interface description.

exclude-channel (Optional) Exclude the channel information from the interface description.

exclude-sub-interface —(Optional) Exclude the subinterface from the interface description.

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 9.1.

`exclude-adapter` and `exclude-sub-interface` options added in Junos OS Release 10.4.

`exclude-channel` option added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

Standard and Vendor-Specific RADIUS Attributes | 3

interface-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 1967
- Hierarchy Level | 1967
- Description | 1968
- Required Privilege Level | 1968
- Release Information | 1968

Syntax

```
interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the interface name is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

interface-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1969
- Hierarchy Level | 1969
- Description | 1969
- Required Privilege Level | 1969

Release Information | 1969

Syntax

```
interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[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 ...]
```

Description

Specify that the interface name is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4

Support at the [edit ... dual-stack-group *dual-stack-group-name*] hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

interface-mib (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 1970
- Hierarchy Level | 1970
- Description | 1970
- Default | 1971
- Required Privilege Level | 1971
- Release Information | 1971

Syntax

```
interface-mib;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name]
```

Description

Enable representation of the Interfaces MIB for the specified dynamic interface.

To achieve maximum performance with enhanced subscriber management, we recommend that you *not* enable representation of the Interfaces MIB on all dynamic subscriber interfaces.

Default

If you do not include the `interface-mib` statement, representation of the Interfaces MIB on dynamic subscriber interfaces is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

Support for the `interface-mib-oids.sh` script added in 18.1R1.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

interface-set (ANCP)

IN THIS SECTION

- Syntax | **1972**
- Hierarchy Level | **1972**
- Description | **1972**
- Options | **1972**
- Required Privilege Level | **1972**
- Release Information | **1972**

Syntax

```
interface-set interface-set-name {
    access-identifier identifier-string;
    underlying-interface underlying-interface-name;
}
```

Hierarchy Level

```
[edit protocols ancp interfaces]
```

Description

Identify a group of VLANs on which traffic is sent to a subscriber identified by the access-loop circuit identifier.

Options

interface-set-name—Name of a group of VLANs that carry traffic to the subscriber identified by the access loop circuit identifier.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Associating an Access Node with Subscribers for ANCP Agent Operations | 1256

interface-traceoptions (DHCP)

IN THIS SECTION

- Syntax | 1973
- Hierarchy Level | 1973
- Description | 1973
- Options | 1974
- Required Privilege Level | 1975
- Release Information | 1975

Syntax

```
interface-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;
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Configure extended DHCP tracing operations that can be enabled on a specific interface or group of interfaces.

Replaces deprecated `interface-traceoptions` statements at the `[edit forwarding-options dhcp-relay]` and `[edit system services dhcp-local-server]` hierarchy levels.

To enable the tracing operation on the specific interfaces, you use the `interface` *`interface-name`* `trace` statement.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements

- `all`—Trace all events
- `packet`—Trace packet and option decoding operations
- `state`—Trace changes in state

`level`—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB),

megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*`k` to specify KB, *size*`m` to specify MB, or *size*`g` to specify GB
- **Range:** 10240 through 1073741824

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Tracing Extended DHCP Operations | 902

interfaces (ANCP)

IN THIS SECTION

- Syntax | 1976
- Hierarchy Level | 1976
- Description | 1976
- Options | 1976
- Required Privilege Level | 1976
- Release Information | 1976

Syntax

```
interfaces {
    interface-set interface-set-name {
        access-identifier identifier-string;
        underlying-interface underlying-interface-name;
    }
    interface-name {
        access-identifier identifier-string
    }
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Identify the subscribers whose traffic is reported and shaped by the ANCP agent.

Options

interface-name—Name of a logical interface supporting a single VLAN that carries traffic to the subscriber identified by the access-loop circuit identifier.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Associating an Access Node with Subscribers for ANCP Agent Operations | 1256

interfaces (Static and Dynamic Subscribers)

IN THIS SECTION

- Syntax | 1977
- Hierarchy Level | 1982
- Description | 1982
- Options | 1983
- Required Privilege Level | 1983
- Release Information | 1983

Syntax

```
interfaces {
    interface-name {
        unit logical-unit-number {
            actual-transit-statistics;
            auto-configure {
                agent-circuit-identifier {
                    dynamic-profile profile-name;
                }
                line-identity {
                    include {
                        accept-no-ids;
                        circuit-id;
                        remote-id;
                    }
                    dynamic-profile profile-name;
                }
            }
```

```
family family {
    access-concentrator name;
    address address;
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
            shared-name  filter-shared-name;
        }
        output filter-name {
            precedence precedence;
            shared-name  filter-shared-name;
        }
    }
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
        mode loose;
    }
    service {
         input {
             service-set service-set-name {
                 service-filter filter-name;
             }
             post-service-filter filter-name;
        }
        output {
             service-set service-set-name {
                 service-filter filter-name;
             }
        }
    }
    service-name-table table-name
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-
```

```
seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name (
                    precedence precedence;
                    shared-name  filter-shared-name;
                }
                output filter-name {
                    precedence precedence;
                    shared-name  filter-shared-name;
                }
            }
            host-prefix-only;
            ppp-options {
                chap;
                pap;
            }
            proxy-arp;
            service {
                pcef  pcef-profile-name {
                    activate rule-name | activate-all;
                }
            }
            targeted-options {
                backup backup;
                group group;
                primary primary;
                weight ($junos-interface-target-weight | weight-value);
            }
            vlan-id;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
        vlan-tagging;
    }
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
```

```
        }
        pppoe-underlying-options {
            max-sessions number;
        }
    }
    demux0 {
        unit logical-unit-number {
            demux-options {
                underlying-interface interface-name
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                demux-source {
                    source-prefix;
                }
                filter {
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
                mac-validate (loose | strict):
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    fail-filter filter-name;
                    mode loose;
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-
seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name;
```

```
                output filter-name;
            }
            vlan-id number;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
    }
    pp0 {
        unit logical-unit-number {
            keepalives interval seconds;
            no-keepalives;
            pppoe-options {
                underlying-interface interface-name;
                server;
            }
            ppp-options {
                aaa-options aaa-options-name;
                authentication [ authentication-protocols ];
                chap {
                    challenge-length minimum minimum-length maximum maximum-length;
                    local-name name;
                }
                ignore-magic-number-mismatch;
                initiate-ncp (dual-stack-passive | ipv6 | ip)
                ipcp-suggest-dns-option;
                mru size;
                mtu (size | use-lower-layer);
                on-demand-ip-address;
                pap;
                peer-ip-address-optional;
                local-authentication {
                    password password;
                    username-include {
                        circuit-id;
                        delimiter character;
                        domain-name name;
                        mac-address;
                        remote-id;
                    }
                }
            }
            family inet {
                unnumbered-address interface-name;
                address address;
```

```
                service {
                    input {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                filter {
                    input filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
            }
        }
    }
    stacked-interface-set {
        interface-set-name interface-set-name {
            interface-set-name interface-set-name;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Define interfaces for dynamic client profiles.

Options

interface-name—The interface variable (`$junos-interface-ifd-name`). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

> **NOTE**: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *DHCP Subscriber Interface Overview*
- Subscribers over Static Interfaces Configuration Overview | **1487**
- Demultiplexing Interface Overview

interim-rate (Access)

IN THIS SECTION

- Syntax | 1984
- Hierarchy Level | 1984
- Description | 1984
- Options | 1984
- Required Privilege Level | 1984
- Release Information | 1985

Syntax

```
interim-rate rate;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the rate at which RADIUS interim update requests are processed.

Options

rate—Maximum number of RADIUS requests sent per second.

- **Range:** 50 through 4000 RADIUS requests per second
- **Default:** 500 RADIUS requests per second

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 12.2.

RELATED DOCUMENTATION

Configuring Options that Apply to All RADIUS Servers | 480

ip-address-first

IN THIS SECTION

- Syntax | 1985
- Hierarchy Level | 1985
- Description | 1986
- Required Privilege Level | 1986
- Release Information | 1986

Syntax

```
ip-address-first;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Configure the extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **773**
- Understanding Differences Between Legacy DHCP and Extended DHCP | **691**
- Address-Assignment Pools Overview | **1139**

ip-can-type (PCRF Partition)

IN THIS SECTION

- Syntax | **1987**
- Hierarchy Level | **1987**
- Description | **1987**
- Options | **1987**
- Required Privilege Level | **1987**
- Release Information | **1987**

Syntax

```
ip-can-type number;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure an IP connectivity access network (IP-CAN) value according to what best fits your operating environment and access network. The Policy and Charging Rules Function (PCRF) partition requires that you configure this statement. The Policy Control and Charging (PCC) enables a centralized control to ensure that the service sessions (IP-CAN sessions) are provided with appropriate bandwidth and QoS.

An IP-CAN bearer is the IP transmission path of defined capacity, delay, and bit error rate. An IP-CAN session incorporates one or more IP-CAN bearers. Support for multiple IP-CAN bearers per IP-CAN session is IP-CAN specific. An IP-CAN session exists as long as the related IPv4 address or IPv6 prefix is assigned and announced to the IP network.

If an IP-CAN session is modified, the Policy and Charging Enforcement Function (PCEF) first uses the event trigger to determine whether to request the PCC rules for the modified IP-CAN session from the PCRF. Then upon reception of updated PCC rules from the PCRF, the PCEF activates, modifies, or removes the PCC rules as indicated by the PCRF.

Options

number Identifier of the IP-CAN value used for your operating environment and access network.

- **Default:** 77

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 16.2.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

jsrc (JSRC)

IN THIS SECTION

- Syntax | 1988
- Hierarchy Level | 1988
- Description | 1989
- Required Privilege Level | 1989
- Release Information | 1989

Syntax

```
jsrc {
    dualstack-support;
    partition partition-name  {
        diameter-instance instance-name;
        destination-host hostname;
        destination-realm realm-name;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure JSRC to interact with an SAE in an SRC environment to authorize and provision subscribers.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

jsrc (Access Profile)

IN THIS SECTION

- Syntax | 1990
- Hierarchy Level | 1990
- Description | 1990
- Required Privilege Level | 1990
- Release Information | 1990

Syntax

```
jsrc {
    attributes {
        exclude {
            user-name [ authorization-request | provisioning-request ];
            }
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Specify JSRC settings in an access profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- JSRC Configuration Overview | **1476**
- Authorizing Subscribers with JSRC | **1478**
- Understanding JSRC-SAE Interactions | **1469**
- Excluding AVPs from Diameter Messages for JSRC | **1480**

jsrc-partition

IN THIS SECTION

- Syntax | 1991
- Hierarchy Level | 1991
- Description | 1991
- Options | 1991
- Required Privilege Level | 1991
- Release Information | 1992

Syntax

```
jsrc-partition partition-name;
```

Hierarchy Level

```
[edit]
```

Description

Specify the JSRC partition to use.

Options

partition-name—Name of the JSRC partition that you want JSRC to use. The name is defined with the `partition` statement at the `[edit jsrc]` hierarchy level.

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 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Configuring the JSRC Partition | 1477

layer2-unicast-replies

IN THIS SECTION

- Syntax | 1992
- Hierarchy Level | 1992
- Description | 1993
- Required Privilege Level | 1993
- Release Information | 1993

Syntax

```
layer2-unicast-replies;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Override the setting of the broadcast bit in DHCP request packets and instead use the Layer 2 unicast transmission method to transmit DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

dhcp-relay | 1752

keep-incoming-circuit-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1994
- Hierarchy Level | 1994

- Description | 1994
- Required Privilege Level | 1994
- Release Information | 1995

Syntax

```
keep-incoming-circuit-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id]
```

Description

Specify that the jdhcpd process keeps the incoming circuit ID and prepends the ID with the locally generated ID (in the format, generated-id + incoming-id) before sending the leasequery packet to the DHCP server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the circuit ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the circuit ID to authenticate subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

keep-incoming-interface-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1995
- Hierarchy Level | 1995
- Description | 1996
- Required Privilege Level | 1996
- Release Information | 1996

Syntax

```
keep-incoming-interface-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
```

Description

Specify that the `jdhcpd` process keeps the incoming interface ID and prepends the ID with the locally generated ID (in the format, `generated-id + incoming-id`) before sending the leasequery packet to the DHCPv6 server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the interface ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the interface ID to authenticate subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | **919**

keep-incoming-remote-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **1997**
- Hierarchy Level | **1997**
- Description | **1997**
- Required Privilege Level | **1997**
- Release Information | **1998**

Syntax

```
keep-incoming-remote-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82 XXXcircuit-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id],
```

Description

Specify that the jdhcpd process keeps the incoming remote ID and prepends the ID with the locally generated ID (in the format, generated-id + incoming-id) before sending the leasequery packet to the DHCPv6 server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the remote ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the remote ID to authenticate subscribers.

Use the statement at the [edit ... dhcpv6] hierarchy level to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | 919

leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1998
- Hierarchy Level | 1998
- Description | 1999
- Options | 1999
- Required Privilege Level | 1999
- Release Information | 2000

Syntax

```
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure parameters that the DHCP relay agent uses when sending DHCP leasequery messages to obtain lease information from the DHCP local servers in the logical system/routing instance.

NOTE: You must also configure support on the relevant DHCP local servers with the "allow-leasequery" on page 1612 statement.

Options

attempts *number-of-attempts*

Specify the number of times the DHCP relay agent attempts to send DHCP leasequery messages to the configured DHCP servers in the logical system/routing instance. DHCP relay agent resends the query message if the configured `timeout` value is reached, and either a confirmed reply or a reply from all configured DHCP servers has not been received. DHCP relay agent sends the subsequent messages only to the DHCP servers that have not replied to previous queries.

- **Range:** 1 through 10
- **Default:** 6

timeout *seconds*

Specify the number of seconds that DHCP relay agent waits before resending a leasequery message to the DHCP servers when all servers have not responded to a previous message.

- **Range:** 1 through 10
- **Default:** 10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Configuring and Using DHCP Individual Leasequery | 812

lease-time-threshold (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2000
- Hierarchy Level | 2000
- Description | 2001
- Options | 2001
- Required Privilege Level | 2001
- Release Information | 2001

Syntax

```
lease-time-threshold seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 lease-time-validation],
[edit forwarding-options dhcp-relay group group-name lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 group group-name lease-time-validation],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
```

```
[edit system services dhcp-local-server lease-time-validation],
[edit system services dhcp-local-server dhcpv6 lease-time-validation],
[edit system services dhcp-local-server group group-name lease-time-validation],
[edit system services dhcp-local-server dhcpv6 group group-name lease-time-validation]
```

Description

Configure the minimum DHCP lease time allowed in your subscriber access network. If a third-party DHCP server or address pool provides a client lease that is less than the configured threshold, the router performs the action specified by the `violation-action` statement.

Options

seconds Minimum client lease time allowed.

- **Range:** 60 through 2,147,483,647 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold | 783

lease-time-validation (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2002
- Hierarchy Level | 2002
- Description | 2003
- Required Privilege Level | 2003
- Release Information | 2003

Syntax

```
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Enable the lease-time validation feature on the router. You can then configure the lease-time threshold and an optional action to take when a lease-time violation occurs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold | 783

limit

IN THIS SECTION

- Syntax | 2004
- Hierarchy Level | 2004
- Description | 2004
- Options | 2004
- Required Privilege Level | 2004
- Release Information | 2004

Syntax

```
limit max-sub-sessions;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name max-data-sessions-
per-subscriber]
```

Description

Specify the limit for the maximum number of subscriber sessions.

Options

max-sub-sessions—Maximum number of subscriber sessions.

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 11.4.

linked-pool-aggregation (Address-Assignment Pools)

IN THIS SECTION

- Syntax | **2005**
- Hierarchy Level | **2005**
- Description | **2005**

Required Privilege Level | 2005
Release Information | 2006

Syntax

```
linked-pool-aggregation;
```

Hierarchy Level

```
[edit access]
```

Description

Change how the search for an available address is performed. When you include this statement, all ranges in the matching pool are searched, but only from the last-saved next address to the highest address in the range. The same search is performed as necessary in a linked pool, through the end of the chain of pools. Then the search begins again at the first pool in the chain; all ranges are searched in each pool, but this time the entire range is searched, from the lowest to the highest address in the range.

This nondefault behavior enables addresses to be assigned non-contiguously, meaning that a free address can be allocated from farther down the chain of linked pools even when free addresses are available in the matching pool or the first pool in the chain. This may be desirable if you configure your RADIUS server to use the IP address alone to identify subscribers.

For example, without this statement, a subscriber can disconnect and that address can be assigned to the next subscriber. The Acct-Start for the second subscriber is sent before the Acct-Stop is sent for the disconnected subscriber. When the Acct-Stop is received, the new subscriber, identified only by the IP address, may 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.

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 18.1R1.

RELATED DOCUMENTATION

- Configuring Address-Assignment Pool Linking | 1153
- Address Allocation from Linked Address Pools | 1141
- Address-Assignment Pools Overview | 1139
- Address-Assignment Pool Configuration Overview | 1146

local (Flat-File Access Profile)

IN THIS SECTION

- Syntax | 2006
- Hierarchy Level | 2006
- Description | 2007
- Options | 2007
- Required Privilege Level | 2007
- Release Information | 2007

Syntax

```
local {
    flat-file-profile profile-name;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure local reporting of service accounting statistics according to the specified flat-file profile, which determines the accounting statistics and nonstatistical parameters that are collected and reported in the local flat file. This configuration collects the running total service statistics per interface family. Because the statistics are maintained in the Routing Engine in a statistics database, they are not affected by a line-card restart, a graceful Routing Engine switchover, or a unified in-service software upgrade (ISSU). The statistics counters are reset when the router reboots.

Collecting service statistics into a local flat file is useful when you do not need to send the information to RADIUS, such as when your are using the information for internal monitoring and need an accurate record for analytics.

NOTE: Starting in Junos OS Release 18.4R1, this statement is no longer supported. If included in a configuration, it is ignored.

Options

profile-name	Name of the flat-file 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 17.1.

This statement is no longer supported as of Junos OS Release 18.4R1

RELATED DOCUMENTATION

Configuring Service Accounting in Local Flat Files

Flat-File Accounting Overview

local-decision (PCRF Partition)

IN THIS SECTION

- Syntax | 2008
- Hierarchy Level | 2008
- Description | 2008
- Options | 2009
- Required Privilege Level | 2010
- Release Information | 2011

Syntax

```
local-decision {
    deny;
    grant;
    reinit-on-failure;
    reinit-on-rar;
    reinit-timeout seconds;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure a non-standard extension that allows the subscriber login to proceed even if the Policy and Charging Rule Function (PCRF) is not available or if the PCRF does not respond within the corresponding logout response timeout. This statement determines how long to wait before using the local decision to log in.

You can also use this statement to configure the router to reinitialize the PCRF session if either of two PCRF errors result when the following sequence occurs:

1. The router sends a CCR-GX-I to the PCRF.
2. The PCRF responds with a CCA-GX-I, but the message is lost.
3. The router retries sending a CCR-GX-I to the PCRF.

Table 91 on page 2009 lists the errors and the resolution.

Table 91: Reinitialization Options for PCRF Errors

Error	Resolution
The PCRF server responds to the router's CCR-GX-I with a CCA-GX-I that contains an unable-to-comply error code (5012) in AVP 268.	Include the `reinit-on-failure` option.
The PCRF responds with a RAR message of any type.	Include the `reinit-on-rar` 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.
2. The router waits for a CCA-GX-T for the length of the `reinit-timeout` period.
3. If the router either receives a CCA-GX-T within the timeout or the timeout expires without receiving a CCA-GX-T, the router regenerates the session ID (conveyed in Diameter AVP 263) and appends a session stamp that consists of the UTC time when the router creates the CCR-GX-I.

The `reinit-on-failure` and `reinit-on-rar` options affect existing subscribers only when the PCRF client (local) state is local-active, local-grant, or started.

Options

deny (Optional) Prevent subscriber logins from occurring when the PCRF is not available or not responding.

grant (Optional) Allow subscriber logins to occur when the PCRF is not available or not responding.

- **Default:** deny

reinit-on-failure (Optional) Reinitialize the PCRF session if PCRF responds with an unable-to-comply error code instead of sending another CCA-GX-I, after the router retries sending a CCR-GX-I.

> **NOTE**: You must also configure the following:
>
> - The `grant` option
> - The `use-session-stamp` option with the "partition" on page 2163 statement

reinit-on-rar (Optional) Reinitialize the PCRF session if the PCRF responds with a RAR of any type after the router retries sending a CCR-GX-I.

> **NOTE**: You must also configure the following:.
>
> - The `grant` option
> - The `use-session-stamp` option with the "partition" on page 2163 statement

reinit-timeout *seconds* (Optional) Sets the number of seconds that the router waits for a CCA-GX-T during a local reinitialization operation. If the timeout expires before the CCA-GX-T is received, then the router performs the same actions as if it had received the CCA-GX-T within the timeout period. The router sends another CCR-GX-I to the PCRF with a new extended session ID (conveyed in Diameter AVP 263). The extended session ID has an appended session stamp that consists of the UTC time when the router creates the CCR-GX-I.

- **Default:** 10
- **Range:** 0 through 90

timeout *seconds* Sets the number of seconds to wait before using the local decision to log in.

- **Default:** 90
- **Range:** 0 through 86,400 seconds (24 hours)

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 16.2.

`reinit-on-failure`, `reinit-on-rar`, and `reinit timeout` options added in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | **1455**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**
- Understanding Upstream and Downstream Messages for the PCRF | **1444**

local-server-group (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | **2011**
- Hierarchy Level | **2012**
- Description | **2012**
- Options | **2012**
- Required Privilege Level | **2012**
- Release Information | **2012**

Syntax

```
local-server-group local-server-group;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-with)],
[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 ...]
```

Description

Forward DHCP client packets to the specified group of DHCP local servers when you use the DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

local-server-group	Name of DHCP local server group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

location (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2013
- Hierarchy Level | 2013
- Description | 2013
- Required Privilege Level | 2014
- Release Information | 2014

Syntax

```
location name;
```

Hierarchy Level

```
[edit forwarding-optionsdhcp-relay relay-option-82 vendor-specific],
[edit forwarding-options dhcp-relaygroup group-name relay-option-82 vendor-specific]
```

Description

Supports the addition of a vendor-specific location in the option 82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The location should be specified as interface, vlan ID, and if applicable, svlan ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The the location is option-data 2 (the hostname is option-data 1). The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information,

and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

location (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | **2014**
- Hierarchy Level | **2014**
- Description | **2015**
- Required Privilege Level | **2015**
- Release Information | **2015**

Syntax

```
location name;
```

Hierarchy Level

```
[edit forwarding-optionsdhcp-relaydhcpv6 relay-option-vendor-specific],
[edit forwarding-optionsdhcp-relaygroup group-name dhcpv6 relay-option-vendor-specific]
```

Description

Supports the addition of a vendor-specific location in the vendor-specific option (17) of DHCPv6 control messages on server-facing interfaces. The location should be specified as interface, vlan ID, and if applicable, svlan ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The the location is option-data 2 (the hostname is option-data 1). The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

logical-interface-unit-range

IN THIS SECTION

- Syntax | **2016**
- Hierarchy Level | **2016**
- Description | **2016**
- Options | **2016**
- Required Privilege Level | **2016**
- Release Information | **2016**

Syntax

```
logical-interface-unit-range (high | low)
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Configure the range from which the unit number is selected for the logical interface service that is created by Extensible Subscriber Services Manager by using an op script. Extensible Subscriber Services Manager assigns the first available unit number in the specified range.

Options

high Upper limit of the logical interface unit range.

- **Range:** 1 through 16,385
- **Default:** 16,385

low Lower limit of the logical interface unit range.

- **Range:** 1 through 16,385
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

logical-system (Diameter Peer)

IN THIS SECTION

- Syntax | 2017
- Hierarchy Level | 2017
- Description | 2017
- Options | 2017
- Required Privilege Level | 2018
- Release Information | 2018

Syntax

```
logical-system  logical-system-name  <routing-instance routing-instance-name > ;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify a logical system and optionally a routing instance for a Diameter peer. Alternatively, you can include the `routing-instance` statement at the `[edit diameter peer peer-name]` hierarchy level to configure only a routing instance.

Options

logical-system-name— Name of the logical system.

- **Default:** Default logical system

`routing-instance` *routing-instance-name*—(Optional) Name of the routing instance.

- **Default:** Master routing instance

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

logical-system (Diameter Transport)

IN THIS SECTION

- Syntax | 2018
- Hierarchy Level | 2019
- Description | 2019
- Options | 2019
- Required Privilege Level | 2019
- Release Information | 2019

Syntax

```
logical-system logical-system-name <routing-instance routing-instance-name >;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Specify a logical system and optionally a routing instance for the transport layer connection.

> **NOTE**: The logical system and routing instance must match that for the peer or a configuration error is reported.

Options

logical-system-name—Name of the logical system.

- **Default:** Default logical system

`routing-instance` *routing-instance-name*—(Optional) Name of the routing instance.

- **Default:** Master routing instance

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 11.2.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring the Diameter Transport | 1376

logical-system-name (Static Subscribers)

IN THIS SECTION

- Syntax | 2020
- Hierarchy Level | 2020
- Description | 2021
- Required Privilege Level | 2021
- Release Information | 2021

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that the name of the logical system is included as part of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487
- Configuring the Static Subscriber Global Username | 1492
- Configuring the Static Subscriber Group Username | 1497

logical-system-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 2022
- Hierarchy Level | 2022
- Description | 2022
- Required Privilege Level | 2022
- Release Information | 2022

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...]
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify that the logical system name be concatenated with the username during the subscriber authentication or DHCP client process. No logical system name is concatenated if the configuration is in the default logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

logical-system-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2023
- Hierarchy Level | 2023
- Description | 2024
- Required Privilege Level | 2024
- Release Information | 2024

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify that the logical system name is concatenated with the username during the subscriber authentication or client authentication process. No logical system name is concatenated if the configuration is in the default logical system. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

logout-response-timeout (PCRF Partition)

IN THIS SECTION

- Syntax | 2025
- Hierarchy Level | 2025
- Description | 2025
- Options | 2026
- Required Privilege Level | 2026
- Release Information | 2026

Syntax

```
logout-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the amount of time in seconds before a Policy and Charging Rule Function (PCRF) partition stops attempting to send a subscriber logout.

If you set `draining` and the `draining-response-timeout` statements for the PCRF partition, any new subscriber logins are denied, and the time limit you set in the `draining-response-timeout` statement is used instead of the `logout-response-timeout` time limit.

NOTE: Any configuration changes made to this statement apply to all subscribers currently waiting to log out within a 60 second period.

Options

seconds Number of seconds to wait before a PCRF partition stops attempting to send a subscriber logout.

- **Default:** 7200
- **Range:** 0 through 86,400 seconds (24 hours)

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | **1455**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | **1418**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

Understanding Upstream and Downstream Messages for the PCRF | **1444**

ltv-syslog-interval (System Process)

IN THIS SECTION

- Syntax | **2027**
- Hierarchy Level | **2027**
- Description | **2027**
- Options | **2027**

Required Privilege Level | 2027
Release Information | 2027

Syntax

```
ltv-syslog-interval seconds;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Configure how often the router logs consolidated syslog messages for DHCP lease-time violations.

Options

seconds Time interval that specifies how often the router logs syslog messages.

- **Range:** 600 through 86,400 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold | 783

processes

mac-address (DHCP Local Server)

IN THIS SECTION

- Syntax | 2028
- Hierarchy Level | 2028
- Description | 2029
- Required Privilege Level | 2029
- Release Information | 2029

Syntax

```
mac-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
```

```
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.

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.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

mac-address (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2030
- Hierarchy Level | 2030
- Description | 2030
- Required Privilege Level | 2031
- Release Information | 2031

Syntax

```
mac-address;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[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 ...]
```

Description

Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or client authentication process.

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.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

maintain-subscriber (Subscriber Management)

IN THIS SECTION

- Syntax | 2032
- Hierarchy Level | 2032
- Description | 2032
- Required Privilege Level | 2032
- Release Information | 2033

Syntax

```
maintain-subscriber {
    interface-delete;
}
```

Hierarchy Level

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

Description

Configure the router to maintain, rather than log out, DHCP relay and DHCP local server based subscribers when the specified type of event occurs.

For example, by default, the router logs out DHCP subscribers when an interface delete event occurs, such as a line card reboot or failure. You would specify the `interface-delete` option to ensure that the router maintains subscribers during line card reboots or failures. However, this option does not maintain subscribers during router reboots or failures.

This statement provides a global configuration for the router, which applies to all DHCP local server and DHCP relay clients in all routing instances.

NOTE: The `maintain-subscriber` statement and `remove-when-no-subscribers` statement are mutually exclusive. You cannot specify that dynamically configured VLAN interfaces are removed when no subscribers exist when the router is also configured to maintain subscribers.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Maintaining Subscribers During Interface Delete Events | 865

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events | 866

managed-configuration (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2033
- Hierarchy Level | 2033
- Description | 2034
- Default | 2034
- Required Privilege Level | 2034
- Release Information | 2034

Syntax

```
(managed-configuration | no-managed-configuration);
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable the dynamic host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured:

- `managed-configuration`—Enable host to use stateful autoconfiguration.
- `no-managed-configuration`—Disable host from using stateful autoconfiguration.

Default

The configured object is disabled unless explicitly enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

map (Domain Map)

IN THIS SECTION

- Syntax | 2035
- Hierarchy Level | 2036
- Description | 2036
- Options | 2036

Required Privilege Level | 2036

Release Information | 2037

Syntax

```
map domain-map-name {
    aaa-logical-system logical-system-name {
        aaa-routing-instance routing-instance-name;
    }
    aaa-routing-instance routing-instance-name;
    access-profile profile-name;
    address-pool pool-name;
    dynamic-profile profile-name;
    strip-domain;
    strip-username (left-to-right | right-to-left);
    sub-domain name {
   (
   aaa-logical-system name {aaa-routing-instance (default | name)
   } | aaa-routing-instance (default | name));
   (
   target-logical-system name {target-routing-instance (default | name)
   } | target-routing-instance (default | name));
   access-profile access-profile;
   address-pool address-pool;
   dynamic-profile dynamic-profile;
   override-chap-password override-chap-password;
   override-password override-password;
   qualifier {
    vlan-id-list [ vlan-id-list ... ];
   }
   strip-domain;
   strip-username (left-to-right | right-to-left);
   tunnel-profile tunnel-profile;
   using-user-password;
  }
    override-password password;
    target-logical-system logical-system-name {
        target-routing-instance routing-instance-name;
```

```
    }
    target-routing-instance routing-instance-name;
    tunnel-profile profile-name;
    tunnel-switch-profile profile-name;
}
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the domain map to use to map options and parameters to subscriber sessions based on the subscriber domain.

Options

domain-map-name—Name of the domain map. The name is the same as the subscriber domain to which it will apply. For example, for the username `user1@example.com`, the domain map name is `example.com`.

- `*` —Use the asterisk wildcard character in the *domain-map-name* to specify a wildcard domain map, which enables mapping based on a partial match (for example, `xyz*northern.example.com`).The router performs the wildcard lookup when there is no exact match for the subscriber domain name. The wildcard can appear anywhere within the domain name string, and can match zero or more characters. The asterisk is the only wildcard character, and only one wildcard is supported in a domain map name. If you include multiple asterisks, the first asterisk is treated as the wildcard character and the others are treated as non-wildcard characters.
- `default`—Use a domain map name of `default` to specify the domain map that the router uses when there is no exact or wildcard match for the domain or realm name in the subscriber username.
- `none`—Use a domain map name of `none` to specify the domain map the router uses when a subscriber username does not have a domain or realm name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`override-password` and `strip-username` options introduced in Junos OS Release 15.1.

wildcard character introduced in Junos OS Release 16.1.

`sub-domain` option introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Configuring a Domain Map | 659

max-advertisement-interval (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2037
- Hierarchy Level | 2038
- Description | 2038
- Options | 2038
- Required Privilege Level | 2038
- Release Information | 2038

Syntax

```
max-advertisement-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Maximum interval between each router advertisement message.

Options

seconds—Maximum interval.

- **Range:** 4 through 1800 seconds
- **Default:** 600 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | 942

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

max-data-sessions-per-subscriber

IN THIS SECTION

- Syntax | 2039
- Hierarchy Level | 2039
- Description | 2039
- Required Privilege Level | 2040
- Release Information | 2040

Syntax

```
max-data-sessions-per-subscriber {
    limit max-sub-sessions;
    exceed-action {
        drop;
        syslog;
    }
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Specify the maximum number of sessions that are concurrently enabled for the named service. The system randomly selects a number of sessions and enables the named service for them. To limit the service's use of resources, other sessions cannot access these named services.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

max-db-size (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 2040
- Hierarchy Level | 2040
- Description | 2041
- Options | 2042
- Required Privilege Level | 2042
- Release Information | 2042

Syntax

```
max-db-size size;
```

Hierarchy Level

```
[edit system configuration-database]
```

Description

Specify the maximum amount of system shared memory that is available for the Junos OS configuration database and the schema database, together. Starting in Junos OS Release 20.1R1, Junos OS uses a single memory map for both the configuration and schema databases.

In Junos OS Release 19.4Rx and lower-numbered releases, this statement specifies the maximum amount of system shared memory that is available for only the Junos OS configuration database

JUNOS OS processes map shared memory into their process space. For example, on MX240 through MX10003 routers, processes can map up to 1GB of shared memory. Enhanced subscriber management processes contend for shared memory with the JUNOS OS configuration database. Shared memory that is not assigned to the configuration database is automatically available to enhanced subscriber management. By default, the configuration database tries to reserve 80 percent of the shared memory map, leaving insufficient space for subscriber management to function.

The majority of configurations require much less than 300MB of mapped space. An appropriate database size enables subscriber management to operate and scale optimally. In some circumstances, you must limit the size of the database to increase the amount of shared memory available to subscriber management. In other circumstances, we recommend that you allow the router to determine the appropriate size and that you do not configure a maximum size.

- For MX5, MX10, MX40, MX80, and MX104 routers, you must always configure the maximum size to be no more than 100MB, regardless of the which Junos OS release is running and regardless of Routing Engine RAM.
- For MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, and MX10003 routers, the decision whether to explicitly configure a maximum size and what that size is, depends on the Junos OS release and the amount of RAM in the Routing Engines:

Junos OS Release	Routing Engine RAM	Recommendation
Release 17.4R1 and earlier releases Release 18.1R1	Any	Configure maximum size to no more than 300MB.
Release 17.4R2 and higher 17.4x releases Release 18.1R2 and higher releases	Routing Engines have at least 32GB each	Allow the router to determine the appropriate size. Do not configure a maximum size.

(Continued)

Junos OS Release	Routing Engine RAM	Recommendation
Release 17.4R2 and higher 17.4x releases Release 18.1R2 and higher releases	Routing Engines have less than 32GB each	Configure maximum size to no more than 300MB.

BEST PRACTICE: You must reboot the device before a change in `max-db-size` value takes effect. By setting this value as part of your initial configuration, you can avoid multiple reboots.

Options

size Specifies the portion of system shared memory, in megabytes (MB), that is allocated for the Junos configuration database.

- **Syntax:** *size* `M` to specify MB
- **Default:** 698,343,424 bytes, only when running Junos OS Release 17.4R1, enhanced IP network services and enhanced subscriber management are enabled, and both Routing Engines in the chassis have at least 32 GB of RAM. Otherwise, there is no default value.

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 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

max-failures

IN THIS SECTION

- Syntax | 2043
- Hierarchy Level | 2043
- Description | 2043
- Options | 2043
- Required Privilege Level | 2044
- Release Information | 2044

Syntax

```
max-failures max-failures;
```

Hierarchy Level

```
[edit logical-systems name protocols ppp-service],
[edit protocols ppp-service]
```

Description

The `max-failures` statement enables you to define the maximum number of failure attempts allowed while establishing the Point-to-Point Protocol (PPP) service.

Options

max-failures Specify the maximum number of failure attempts allowed while establishing the PPP service.

- **Range:** 1 through 10
- **Default:** 5

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

max-outstanding-requests (Diameter Applications)

IN THIS SECTION

- Syntax | **2044**
- Hierarchy Level | **2045**
- Description | **2045**
- Options | **2045**
- Required Privilege Level | **2045**
- Release Information | **2045**

Syntax

```
max-outstanding-requests number;
```

Hierarchy Level

```
[edit access gx-plus global],
[edit access ocs partition partition-name],
[edit access pcrf partition partition-name]
```

Description

Limit the number of outstanding requests that the Diameter-based application (function) can retry to a remote server when the requests are improperly answered. Too many requests risks overloading the server and increases the chance of losing messages.

The `gx-plus` statement limits retries from the Gx-Plus function to the Gx-Plus server using the Gx and JSRC protocols. The `ocs` statement limits retries from the OCS function to the OCS server using the Gy protocol. The `pcrf` statement limits retries from the PCRF function to the PCRF server using the Gx protocol.

Options

number—Number of outstanding requests from the function to the server that can exist at any time.

- **Default:** 40
- **Range:** 2 through 40

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 11.2.

Support for OCS and PCRF partitions introduced in Junos OS Release 16.2 for MX Series routers.

RELATED DOCUMENTATION

Configuring Gx-Plus Global Attributes | **1406**

Configuring Gx-Plus | **1404**

Configuring the OCS Partition | 1449
Configuring the PCRF Partition | 1455
Understanding Gx Interactions Between the Router and the PCRF | 1418
Understanding Gy Interactions Between the Router and the OCS | 1431

max-pending-accounting-stops (Access Profile)

IN THIS SECTION

- Syntax | 2046
- Hierarchy Level | 2046
- Description | 2046
- Options | 2047
- Required Privilege Level | 2047
- Release Information | 2047

Syntax

```
max-pending-accounting-stops number;
```

Hierarchy Level

```
[edit access accounting-backup-options]
```

Description

Set the maximum number of pending accounting stop requests that the router backs up while all the RADIUS accounting servers in the profile are offline.

Options

`number` Number of stops to hold.

- **Range:** 1 through 168,000
- **Default:** 168,000

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 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting | 591

max-withhold-time (Access Profile)

IN THIS SECTION

- Syntax | 2048
- Hierarchy Level | 2048
- Description | 2048
- Options | 2048
- Required Privilege Level | 2048
- Release Information | 2048

Syntax

```
max-withhold-time hold-time;
```

Hierarchy Level

```
[edit access accounting-backup-options]
```

Description

Set the timer that determines how long the router holds pending accounting stop requests. Any remaining accounting stop messages are flushed when the timer expires, even if the accounting server is again online.

Options

hold-time Number of minutes.

- **Range:** 1 through 1440
- **Default:** 60

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 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting | 591

maximum-discovery-table-entries

IN THIS SECTION

- Syntax | 2049
- Hierarchy Level | 2049
- Description | 2049
- Default | 2049
- Options | 2050
- Required Privilege Level | 2050
- Release Information | 2050

Syntax

```
maximum-discovery-table-entries entry-number;
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Specify the maximum number of discovery table entries accepted from all ANCP neighbors or from a particular ANCP neighbor. The number of entries configured for an individual neighbor supersedes the global value. The neighbor can continue to update previously created entries when the maximum has been exceeded, but no new entries are accepted.

Default

No limit on the number of table entries

Options

entry-number—Maximum number of discovery table entries.

- **Range:** 1 through 100,000
- **Default:** 100,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the ANCP Agent | **1253**

Configuring ANCP Neighbors | **1255**

maximum-helper-restart-time

IN THIS SECTION

- Syntax | **2051**
- Hierarchy Level | **2051**
- Description | **2051**
- Options | **2051**
- Required Privilege Level | **2051**
- Release Information | **2051**

Syntax

```
maximum-helper-restart-time seconds;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Specify how long other router processes wait for the ANCP agent to restart before considering it to be down.

Options

seconds—Number of seconds other processes wait for ANCP to restart.

- **Range:** 45 through 600 seconds
- **Default:** 45 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | **1253**

Specifying How Long Processes Wait for the ANCP Agent Restart to Complete | **1259**

maximum-subscribers

IN THIS SECTION

- Syntax | 2052
- Hierarchy Level | 2052
- Description | 2052
- Options | 2052
- Required Privilege Level | 2052
- Release Information | 2053

Syntax

```
maximum-subscribers limit;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Configure the maximum number of subscriber sessions supported at a time.

Options

limit Maximum number of subscribers.

- **Range:** 1 through 2000
- **Default:** 1000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show subscribers | 3057

show subscribers summary | 3110

metric (Diameter Route)

IN THIS SECTION

- Syntax | 2053
- Hierarchy Level | 2053
- Description | 2054
- Options | 2054
- Required Privilege Level | 2054
- Release Information | 2054

Syntax

```
metric route-metric;
```

Hierarchy Level

```
[edit diameter network-statement element-name forwarding route dne-route-name]
```

Description

Specify the metric associated with a destination and function. Together, these three elements define a route reachable through a Diameter network element. A lower metric makes a route more preferred.

Options

route-metric—Metric assigned to the route.

- **Range:** 0 through 255

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

min-advertisement-interval (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2055
- Hierarchy Level | 2055
- Description | 2055
- Options | 2055
- Required Privilege Level | 2055
- Release Information | 2055

Syntax

```
min-advertisement-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Minimum interval between each router advertisement message.

Options

seconds—Minimum interval.

- **Range:** 3 seconds through three-quarter times the maximum advertisement interval value
- **Default:** One-third the maximum advertisement interval value

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

multi-address-embedded-option-response (DHCP Local Server)

IN THIS SECTION

- Syntax | 2056
- Hierarchy Level | 2056
- Description | 2056
- Default | 2057
- Required Privilege Level | 2057
- Release Information | 2057

Syntax

```
multi-address-embedded-option-response;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services system services dhcp-local-server
dhcpv6 ...],
[edit routing-instances routing-instance-name system services system services dhcp-local-server
dhcpv6 ...]
```

Description

Configure DHCPv6 local server to return the DNS server address (DHCPv6 attribute 23) as a suboption in the respective IA_NA or IA_PD headers.

Default

DHCPv6 local server returns the DNS server address as a global DHCPv6 option.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.3 and later releases. (Not supported in Junos OS Release 13.1 and Release 13.2.)

RELATED DOCUMENTATION

Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment | **1169**

Overriding the Default DHCP Local Server Configuration Settings | **707**

nas-port-extended-format

IN THIS SECTION

- Syntax | **2058**
- Hierarchy Level | **2058**
- Description | **2058**
- Options | **2059**
- Required Privilege Level | **2059**
- Release Information | **2059**

Syntax

```
nas-port-extended-format {
    adapter-width bits;
    ae-width bits;
    atm {
        adapter-width bits;
        port-width bits;
        slot-width bits;
        vci-width bits;
        vpi-width bits;
    }
    port-width bits;
    pw-width bits;
    slot-width bits;
    stacked-vlan-width bits;
    vlan-width bits;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width in bits of the fields in the NAS-Port 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 specifies the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, aggregated, Ethernet, VLAN, and S-VLAN.

NOTE: The combined total of the widths of all fields for a subscriber 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.

Options

`adapter-width` *width*—Number of bits in the adapter field.

`ae-width` *width*—(Ethernet subscribers only) Number of bits in the aggregated Ethernet identifier field.

`atm`—Specify width for fields for ATM subscribers.

`port-width` *width*—Number of bits in the port field.

`pw-width` *width*—(Ethernet subscribers only) Number of bits in the pseudowire field. Appears in the Cisco NAS-Port-Info AVP (100).

`slot-width` *width*—Number of bits in the slot field.

`stacked-vlan-width` *width*—Number of bits in the SVLAN ID field.

`vci-width` *width*—(ATM subscribers only) Number of bits in the ATM virtual circuit identifier (VCI) field.

`vlan-width` *width*—Number of bits in the VLAN ID field.

`vpi-width` *width*—(ATM subscribers only) Number of bits in the ATM virtual path identifier (VPI) field.

NOTE: The total of the widths must not exceed 32 bits, or the configuration will fail.

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 9.1.

`ae-width` option added in Junos OS Release 12.1.

`atm` option added in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

`atm` option supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

`pw-width` option added in Junos OS Release 15.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

nas-port-extended-format (Interfaces)

IN THIS SECTION

- Syntax | 2060
- Hierarchy Level | 2060
- Description | 2061
- Options | 2061
- Required Privilege Level | 2061
- Release Information | 2061

Syntax

```
nas-port-extended-format {
    adapter-width bits;
    ae-width bits;
    port-width bits;
    slot-width bits;
    stacked;
    stacked-vlan-width bits;
    vci-width bits:
    vlan-width bits;
    vpi-width bits;
}
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width in bits of the fields in the NAS-Port attribute.

Options

`adapter-width` *`width`*—Number of bits in the adapter field.

`ae-width` *`width`*—Number of bits in the aggregated Ethernet identifier field.

`port-width` *`width`*—Number of bits in the port field.

`slot-width` *`width`*—Number of bits in the slot field.

`stacked`—Include stacked VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format.

`stacked-vlan-width` *`width`*—Number of bits in the SVLAN ID field.

`vci-width` *`width`*—Number of bits in the ATM virtual circuit identifier (VCI) field.

`vlan-width` *`width`*—Number of bits in the VLAN ID field.

`vpi-width` *`width`*—Number of bits in the ATM virtual path identifier (VPI) field.

NOTE: Each field can be 0 through 32 bits wide; however, the 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.

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 12.3.

Options `vci-width` and `vpi-width` introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

Options `vci-width` and `vpi-width` supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 527

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 526

nas-port-id-format (Subscriber Management)

IN THIS SECTION

- Syntax | 2062
- Hierarchy Level | 2063
- Description | 2063
- Default | 2063
- Options | 2064
- Required Privilege Level | 2064
- Release Information | 2064

Syntax

```
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    concatenated-vlan-tags (fixed-size-inner-tag | fixed-size-outer-tag)
    interface-description;
    interface-text-description;
    nas-identifier;
    order (agent-circuit-id | agent-remote-id | interface-description | interface-text-
description | nas-identifier | postpend-vlan-tags);
    postpend-vlan-tags;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the optional information that the router includes in the NAS-Port-ID (RADIUS attribute 87) that is passed to the RADIUS server during authentication and accounting. You can include any combination of the optional values.

When you specify the values for the NAS-Port-ID, you can configure the values to appear in either the default order or a custom order of your choice.

NOTE: The default and custom order methods are mutually exclusive. The configuration fails if you attempt to configure a NAS-Port-ID that includes values in both types of orders.

To specify that the optional values appear in the default order in the NAS-Port-ID, configure the values directly under the `nas-port-id-format` statement. The default order is as follows, in which the `#` character is the delimiter:

`nas-identifier # interface-description # interface-text-description # agent-circuit-id # agent-remote-id # postpend-vlan-tags`

To specify a custom order for the NAS-Port-ID string, you use the `order` option. Include the `order` option before each optional value you want to include in the string, in the order in which you want the options to appear. For example, the configuration, `order interface-text-description order nas-identifier order agent-remote-id` produces the following NAS-Port-ID, in which the `#` character is the delimiter:

`interface-text-description # nas-identifier # agent-remote-id`

Starting in Junos OS Release 21.3R1, we have introduced a new NAS-Port-ID format for RADIUS server access request. The NAS-Port-ID format is `S-VLAN<concatenated 0's>C-VLAN:S-VLAN-C-VLAN`. For customer-VLAN (C-VLAN), if the number of digits is less than four, prepend it with zeroes. For example, NAS-Port-ID for an S-VLAN 72 and C-VLAN 82 is `720082:72-82`.

Default

The router includes the interface description in the NAS-Port-ID when no optional values are specified.

Options

agent-circuit-id—Include the agent circuit ID from either DHCP option 82 or the DSL forum VSAs.

agent-remote-id—Include the agent remote ID from either DHCP option 82 or the DSL forum VSAs.

concatenated-vlan-tags—Include the vlan tags as a concatenated string.

fixed-size-inner-tag—Fixed size inner VLAN tag value of 4 octets.

fixed-size-outer-tag—Fixed size outer VLAN tag value of 4 octets.

interface-description—Include the interface description (interface identifier).

interface-text-description—Include the textual interface description (the text description that is statically configured in the CLI).

nas-identifier—Include the NAS identifier value (RADIUS attribute 32).

order—Specify the optional values you want to include in the NAS-Port-ID and the customized order in which you want the values to appear. You must include the order option before each optional value (for example, order agent-circuit-id order interface-description).

postpend-vlan-tags—Include the VLAN tags. The router includes the tags in the format :<outer-tag>-<inner-tag> for a double-tagged VLAN, or :<outer-tag> for a single-tagged VLAN.

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 11.4.

Options interface-text-description, order, and postpend-vlan-tags introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers | 483

Configuring a NAS-Port-ID with Additional Options | 520

RADIUS Servers and Parameters for Subscriber Access | 476

nas-port-options (RADIUS Options)

IN THIS SECTION

- Syntax | 2065
- Hierarchy Level | 2065
- Description | 2066
- Options | 2066
- Required Privilege Level | 2066
- Release Information | 2066

Syntax

```
nas-port-options nas-port-options-name {
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked;
        stacked-vlan-width width;
        vci-width width:
        vlan-width width;
        vpi-width width;
    }
    nas-port-type port-type;
    stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
    vlan-ranges (any | low-tag-high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name radius-options]
```

Description

Create a NAS-Port options definition to 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. Each NAS-Port options definition includes the NAS-Port extended format, the NAS-Port-Type, and either the VLAN range of subscribers or the S-VLAN range of subscribers to which the definition applies.

NOTE: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

nas-port-options-name	Name of the NAS-Port options definition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **527**

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **526**

nas-port-type (Subscriber Management)

IN THIS SECTION

- Syntax | 2067
- Hierarchy Level | 2067
- Description | 2067
- Default | 2068
- Options | 2068
- Required Privilege Level | 2069
- Release Information | 2069

Syntax

```
nas-port-type {
    ethernet {
        port-type;
    }
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the port type used to authenticate subscribers. The router includes the port type in RADIUS attribute 61 (NAS-Port-Type attribute).

NOTE: This statement is ignored if the `ethernet-port-type-virtual` statement is included in the same access profile.

Default

The router uses a port type of `ethernet`.

Options

`port-type`—One of the following port types:

- *value*—A value from 0-65535
- `adsl-cap`—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- `adsl-dmt`— Asymmetric DSL, discrete multitone (DMT)
- `async`—Asynchronous
- `cable`—Cable
- `ethernet`—Ethernet
- `fddi`—Fiber Distributed Data Interface
- `g3-fax`—G.3 Fax
- `hdlc-clear-channel`—HDLC Clear Channel
- `iapp`—Inter-Access Point Protocol (IAPP)
- `idsl`—ISDN DSL
- `isdn-sync`—ISDN Synchronous
- `isdn-v110`—ISDN Async V.110
- `isdn-v120`—ISDN Async V.120
- `piafs`—Personal Handyphone System (PHS) Internet Access Forum Standard
- `sdsl`—Symmetric DSL
- `sync`—Synchronous
- `token-ring`—Token Ring
- `virtual`—Virtual
- `wireless`—Other wireless
- `wireless-1x-ev`—Wireless 1xEV

- `wireless-cdma2000`—Wireless code division multiple access (CDMA) 2000
- `wireless-ieee80211`—Wireless 802.11
- `wireless-umts`—Wireless universal mobile telecommunications system (UMTS)
- `x25`—X.25
- `x75`—X.75
- `xdsl`—DSL of unknown type

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 11.4.

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers | 483

RADIUS Servers and Parameters for Subscriber Access | 476

nas-port-type (RADIUS Options)

IN THIS SECTION

- Syntax | 2070
- Hierarchy Level | 2070
- Description | 2070
- Default | 2070
- Options | 2070
- Required Privilege Level | 2071

Release Information | 2072

Syntax

```
nas-port-type port-type;
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Specify the port type used to authenticate subscribers. The router includes the port type in the NAS-Port-Type (61) RADIUS IETF attribute.

Default

If you do not include the `nas-port-type` statement at the `[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]` hierarchy level, the global value configured for `nas-port-type` at the `[edit access profile profile-name radius options]` hierarchy level takes effect.

Options

port-type One of the following port types:

- *value*—A value from 0 through 65535
- `adsl-cap`—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- `adsl-dmt`— Asymmetric DSL, discrete mutilating (DOT)
- `async`—Asynchronous
- `cable`—Cable
- `ethernet`—Ethernet

- `fddi`—Fiber Distributed Data Interface
- `g3-fax`—G.3 Fax
- `hdlc-clear-channel`—HDLC Clear Channel
- `iapp`—Inter-Access Point Protocol (IAPP)
- `idsl`—ISDN DSL
- `isdn-sync`—ISDN Synchronous
- `isdn-v110`—ISDN Async V.110
- `isdn-v120`—ISDN Async V.120
- `piafs`—Personal Handyphone System (PHS) Internet Access Forum Standard
- `sdsl`—Symmetric DSL
- `sync`—Synchronous
- `token-ring`—Token Ring
- `virtual`—Virtual
- `wireless`—Other wireless
- `wireless-1x-ev`—Wireless 1xEV
- `wireless-cdma2000`—Wireless code division multiple access (CDMA) 2000
- `wireless-ieee80211`—Wireless 802.11
- `wireless-umts`—Wireless universal mobile telecommunications system (UMTS)
- `x25`—X.25
- `x75`—X.75
- `xdsl`—DSL of unknown type

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 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 527

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 526

nasreq (Diameter Application)

IN THIS SECTION

- Syntax | 2072
- Hierarchy Level | 2073
- Description | 2073
- Options | 2073
- Required Privilege Level | 2074
- Release Information | 2074

Syntax

```
nasreq
    max-outstanding-requests number;
    partition partition-name {
        destination-host hostname;
        destination-realm realm-name;
        diameter-instance master;
    }
    request-retry retries;
```

```
    timeout seconds;
}
```

Hierarchy Level

```
[edit access]
```

Description

Specify the destination and transmission parameters for the Diameter Network Access Server Requirements (NASREQ) protocol.

NASREQ is a Diameter-based authentication and authorization protocol. The NASREQ client has two queues, the transmit queue and response queue. The transmit queue stores outbound packets until sent to Diameter, and includes requests and responses. The response queue stores request packets until Diameter responds to the request, and includes only requests waiting for a response.

The following configuration options control transmission flow and use of the queues:

- `max-outstanding-requests` option specifies the maximum number of requests (includes AAR and STR) that the NASREQ client sends 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.
- The `request-retry` option specifies how many times the NASREQ client retries transmitting a packet to the Diameter server when a timeout is received from the Diameter server for the request.
- The `timeout` option specifies the number of seconds that an outbound packet remains in the transmit queue before it is declared to have timed out. The NASREQ client does not transmit packets that have timed out. The timeout value applies to all packets in the transmit queue, including both requests and responses to be sent. Diameter manages packets that time out after transmission.

Options

number Maximum number of requests that the NASREQ client sends to Diameter.

- **Range:** 20 through 100

retries Number of times to re-send a timeout request to Diameter after the initial request. This value applies only to requests in the response queue. A value of 0 indicates to try transmission once and then do not retry.

- **Range:** 0 through 3

seconds Number of seconds that an outbound packet remains in the transmit queue before it is declared to have timed out.

- **Range:** 5 through 30

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Diameter Network Access Server Application (NASREQ) | **1463**

Configuring the Diameter Network Access Server Application (NASREQ) | **1465**

neighbor (Define ANCP)

IN THIS SECTION

- Syntax | **2075**
- Hierarchy Level | **2075**
- Description | **2075**
- Options | **2075**
- Required Privilege Level | **2075**
- Release Information | **2075**

Syntax

```
neighbor ip-address {
    adjacency-loss-hold-time seconds;
    adjacency-timer;
    auto-configure-trigger interface interface-name;
    ietf-mode;
    maximum-discovery-table-entries entry-number;
    pre-ietf-mode;
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an ANCP neighbor with which the ANCP agent on the router forms an adjacency for reporting and shaping traffic.

Options

ip-address—IP address of the ANCP neighbor.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Configuring ANCP Neighbors | 1255

network

IN THIS SECTION

- Syntax | 2076
- Hierarchy Level | 2076
- Description | 2076
- Options | 2076
- Required Privilege Level | 2077
- Release Information | 2077

Syntax

```
network ip-prefix</prefix-length>;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet]
```

Description

Configure subnet information for an IPv4 address-assignment pool.

Options

ip-prefix—IP version 4 address or prefix value.

prefix-length—(Optional) Subnet mask.

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 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

network-element (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 2077
- Hierarchy Level | 2078
- Description | 2078
- Options | 2078
- Required Privilege Level | 2078
- Release Information | 2078

Syntax

```
network-element element-name {
    dne-origin realm realm-name <host hostname>
    forwarding {
        route dne-route-name {
            destination realm realm-name <host hostname> ;
            function [function-name] <partition partition-name>;
            metric route-metric;
```

```
        }
    }
    function function-name;
    peer peer-name {
        priority priority-number;
    }
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Specify the transport layer Diameter configuration. The Diameter network element includes a list of routes reachable through the Diameter instance, associated functions, and prioritized Diameter peers.

Options

element-name—Name of the network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

network-services

IN THIS SECTION

- Syntax | 2079
- Hierarchy Level | 2079
- Description | 2079
- Default | 2079
- Options | 2080
- Required Privilege Level | 2080
- Release Information | 2080

Syntax

```
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip | lan);
```

Hierarchy Level

```
[edit chassis]
```

Description

Set the router's network services to a specific mode of operation. On MX240, MX480, and MX960 routers, MPC5E and MPC7E power on only if the network services mode configured is `enhanced-ip` or `enhanced-ethernet`.

MX2010 and MX2020 support only `enhanced-ip` and `enhanced-ethernet` network services modes.

Default

- MX80, MX104, MX2010, MX2020—`enhanced-ip`
- MX240, MX480, MX960—`ip`

Options

ethernet—Set the router's network services to Ethernet and use standard, compiled firewall filter format.

enhanced-ethernet—Set the router's network services to enhanced Ethernet and use enhanced mode capabilities. Only MPCs and MS-DPCs are powered on in the chassis.

ip—Set the router's network services to Internet Protocol and use standard, compiled firewall filter format.

enhanced-ip—Set the router's network services to enhanced Internet Protocol and use enhanced mode capabilities. Only MPCs and MS-DPCs are powered on in the chassis. Non-service DPCs do not work with enhanced network services mode options. This feature is enabled by default on MX80, MX104, MX2010, and MX2020 Universal Routing Platforms. For MX960 platform, enhanced-ip configuration must be enabled on both REs together. After committing the configuration, GRES configuration must be disabled and both REs must be rebooted to ensure that all the FPCs reboot and have the same network-services as REs. Any mismatch in network services between RE0, RE1, FPCs will lead to unexpected results.

lan—Set the router's network services to LAN and use standard, compiled firewall filter format. Reboot the system after setting the router's network services to LAN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.5.

enhanced-ethernet and enhanced-ip options introduced in Junos OS Release 11.4.

limited-ifl-scaling option introduced in Junos OS Release 15.1R3 for MX Series routers.

RELATED DOCUMENTATION

- Network Services Mode Overview
- Firewall Filters and Enhanced Network Services Mode Overview
- Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers
- Configuring Enhanced IP Network Services for a Virtual Chassis

Limiting the Maximum Number of Logical Interfaces on MX Series Routers With MS-DPCs in Enhanced IP Network Services Mode

no-bind-on-request (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2081
- Hierarchy Level | 2081
- Description | 2082
- Required Privilege Level | 2082
- Release Information | 2082

Syntax

```
no-bind-on-request;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Explicitly disable automatic binding of received DHCP request messages that have no entry in the database (*stray* requests). Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

NOTE: Beginning with Junos OS Release 10.4, automatic binding of stray requests is enabled by default. In Junos OS Release 10.3 and earlier releases, automatic binding of stray requests is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Overriding the Default DHCP Relay Configuration Settings | 709
Disabling Automatic Binding of Stray DHCP Requests | 714

no-unsolicited-ra (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 2083
- Hierarchy Level | 2083
- Description | 2083
- Required Privilege Level | 2083
- Release Information | 2084

Syntax

```
no-unsolicited-ra;
```

Hierarchy Level

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

Description

Disable the default transmission and periodic refresh of unsolicited Router Advertisement messages by the router when the subscriber interface is created, and at configured periodic intervals thereafter. When you include the `no-unsolicited-ra` statement, the router sends Router Advertisement messages and associated periodic refresh messages only when it receives a Router Solicitation message from the subscriber.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

no-vlan-interface-name

IN THIS SECTION

- Syntax | 2084
- Hierarchy Level | 2084
- Description | 2085
- Required Privilege Level | 2086
- Release Information | 2086

Syntax

```
no-vlan-interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-
```

```
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-
interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

When you do not want bridge domain or VLAN tag information, do not include the VLAN ID nor the VLAN interface name (the default) in the circuit or remote ID value in the DHCP option 82 information.

NOTE: The `no-vlan-interface-name` statement is mutually exclusive with the `use-interface-description` and `use-vlan-id` statements.

When you configure the `no-vlan-interface-name` statement only, the format displays only the Layer 3 interface:

```
irb.subunit
```

NOTE: The *subunit* is required and used to differentiate the interface for remote systems.

When you configure the `no-vlan-interface-name` and `use-interface-description` statements, the format displays the IRB interface description:

```
irb_descr
```

If you configure the `no-vlan-interface-name` and `use-interface-description` statements, and no description for the IRB interface is found, the format displays the IRB interface name:

```
irb.subunit
```

When you configure the no-vlan-interface-name and include-irb-and-l2 statements, the format displays the Layer 2 logical interface name and the IRB interface name:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

When you configure the no-vlan-interface-name, include-irb-and-l2 and use-interface-name statements, the format displays the Layer 2 interface description and the IRB interface name:

```
l2_descr+irb.subunit
```

If you configure the no-vlan-interface-name, include-irb-and-l2 and use-interface-name statements, and no description for the Layer 2 interface is found, the format displays the Layer 2 logical interface name and the IRB interface name:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

NOTE: The EX Series switches that support the no-vlan-interface-name statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

Using DHCP Relay Agent Option 82 Information | 751

Configuring DHCPv6 Relay Agent Options | 917

not-present (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2087
- Hierarchy Level | 2087
- Description | 2090
- Required Privilege Level | 2090
- Release Information | 2090

Syntax

```
not-present {
    drop drop;
    forward-only forward-only;
    local-server-group local-server-group;
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
```

```
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-77],
```

```
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
```

```
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Option not present action

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

ocs (Diameter Applications)

IN THIS SECTION

- Syntax | 2091
- Hierarchy Level | 2092
- Description | 2092
- Required Privilege Level | 2092
- Release Information | 2093

Syntax

```
ocs {
    global {
        service-context-id service-context;
    }
    partition partition-name  {
        called-station-id station-name;
        charging-id number;
        destination-host ocs-hostname;
        destination-realm ocs-realm-name;
        diameter-instance;
        draining;
        draining-response-timeout seconds;
        final-response-timeout seconds;
        force-continue;
        ggsn-address address;
        ggsn-mcc-mncggsn-mcc-mnc;
        max-outstanding-requests number;
        send-origin-state-id number;
        user-name-include  {
            base-interface-name;
            delimiter delimiter-character;
            domain-name my-domain;
            interface-name;
            mac-address;
            nas-port-id;
            origin-host;
            origin-realm;
            user-name;
            user-prefix pref;
        }
            backup {
                limit;
                timeout seconds;
                overflow { deny-login | drop-oldest };
            }
            sftp-backup {
                accumulation-timeout seconds;
                accumulation-count messages;
                accumulation-size bytes;
                logical-system;
```

```
                retry-interval seconds;
                response-timeout seconds;
                routing-instance;
                address ipv4 | ipv6;
                port number;
                directory filename;
                file-name-prefix filename-prefix;
                node-ipv4-address;
                node-ipv6-address;
                ssh-login;
                ssh-connection-linger;
                ssh-log-verbose;
                ssh-passphrase;
            }
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Online Charging System (OCS) global attributes and partition. The OCS interacts with the Policy and Charging Enforcement Function (PCEF). The PCEF optionally reports usage and receives additional authorizations from the OCS using the 3rd Generation Partnership Project (3GPP) Gy protocol. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

on-demand-ip-address

IN THIS SECTION

- Syntax | 2093
- Hierarchy Level | 2093
- Description | 2094
- Default | 2094
- Required Privilege Level | 2094
- Release Information | 2094

Syntax

```
on-demand-ip-address;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit”].
```

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit unit-number ppp-options], [edit protocols ppp-service]
```

Description

For IPv4 and IPv6 dual-stack PPP subscribers, enables on-demand allocation and de-allocation of an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address.

Configuration changes take effect as follows:

- When you change this setting for a dynamic PPP interface (at the `[edit dynamic-profiles]` hierarchy level), the change takes effect only for new subscriber logins.
- When you change this setting for a static PPP interface (at the `[edit interfaces pp0]` hierarchy level, the subscribers on the interface are logged out.
- When you change this setting globally (at the `[edit protocols ppp-service]` hierarchy level), the change takes effect only for new subscriber logins.

If you enable on-demand allocation at both the interface and global levels, the global configuration takes precedence and changes take effect for new subscriber logins.

Default

This functionality is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation | **1115**

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | **1122**

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1122
Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers | 1123

on-demand-address-allocation

IN THIS SECTION

- Syntax | 2095
- Hierarchy Level | 2095
- Description | 2095
- Required Privilege Level | 2096
- Release Information | 2096

Syntax

```
on-demand-address-allocation;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-group
dual-stack-group-name],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name]
```

Description

Enables an address or prefix to be allocated on demand when a dual-stack subscriber session is established. Starting in Junos OS Release 18.1R1, when reauthentication is configured, enables per-

family address allocation as each family's DHCP session is established. In earlier released, applies only to the second family of the dual stack.

> **NOTE**: You must configure `on-demand-address-allocation` if you also configure reauthentication for dual-stack, single-session DHCP subscribers. This is true whether you enable reauthentication with the `reauthenticate` statement or the Reauthenticate-On-Renew VSA (26-206).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- Configuring RADIUS Reauthentication for DHCP Subscribers | **567**
- RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | **556**
- Single-Session DHCP Dual-Stack Overview | **1004**
- Configuring Access Profile Options for Interactions with RADIUS Servers | **483**
- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | **703**
- *Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

on-link (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | **2097**

- Hierarchy Level | 2097
- Description | 2097
- Default | 2097
- Required Privilege Level | 2097
- Release Information | 2098

Syntax

```
(on-link | no-on-link);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-
nameprefix prefix]
```

Description

Specify whether to enable prefixes to be used for onlink determination:

- `no-on-link`—Disable prefixes from being used for onlink determination.
- `on-link`—Enable prefixes to be used for onlink determination.

Default

The configured object is enabled unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

option-order (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2098
- Hierarchy Level | 2098
- Description | 2099
- Options | 2099
- Required Privilege Level | 2100
- Release Information | 2100

Syntax

```
option-order name;
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
```

```
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Options precedence order

Options

name Option number

- Values:
  - `60`—Option 60
  - `77`—Option 77

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

option-15 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2100
- Hierarchy Level | 2101
- Description | 2102
- Required Privilege Level | 2102
- Release Information | 2102

Syntax

```
option-15 {
    default-action {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
```

```
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name
```

```
relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
group name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 relay-option]
```

Description

Specify that the payload of the Option 15 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

option-16 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2103
- Hierarchy Level | 2104
- Description | 2105
- Required Privilege Level | 2105
- Release Information | 2105

Syntax

```
option-16 {
    default-action {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
```

```
            drop drop;
            forward-only forward-only;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
group name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
```

```
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 relay-option]
```

Description

Specify that the payload of the Option 16 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

option-60 (DHCP Local Server)

IN THIS SECTION

- Syntax | 2106
- Hierarchy Level | 2106
- Description | 2106
- Required Privilege Level | 2106
- Release Information | 2106

Syntax

```
option-60;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the payload of Option 60 (Vendor Class Identifier) from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

option-60 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2107
- Hierarchy Level | 2108
- Description | 2109
- Required Privilege Level | 2109
- Release Information | 2109

Syntax

```
option-60 {
    default-action {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
    not-present {
        drop drop;
```

```
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
```

```
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Specify that the payload of the Option 60 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 17.4R1 for MX Series.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

option-77 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2110
- Hierarchy Level | 2111
- Description | 2112
- Required Privilege Level | 2112
- Release Information | 2112

Syntax

```
option-77 {
    default-action {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
```

```
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
```

```
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Specify that the payload of the Option 77 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

option-82 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2113
- Hierarchy Level | 2113
- Description | 2113
- Options | 2113
- Required Privilege Level | 2114
- Release Information | 2114

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name authentication username-include]
```

Description

Specify the option 82 that is concatenated with the username during the subscriber authentication or client authentication process. You can specify either, both, or neither the Agent Circuit ID and the Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If neither suboption is supplied, the raw payload of option 82 is concatenated to the username.

NOTE: The option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.

Options

`circuit-id`—(Optional) The string for the Agent Circuit ID suboption (suboption 1).

`remote-id`—(Optional) The string for the Agent Remote ID suboption (suboption 2).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

option-82 (DHCP Local Server Authentication)

IN THIS SECTION

- Syntax | 2114
- Hierarchy Level | 2115
- Description | 2115
- Options | 2115
- Required Privilege Level | 2115
- Release Information | 2116

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the type of Option 82 information from the client PDU that is concatenated with the username during the subscriber authentication or DHCP client authentication process. You can specify either, both, or neither of the Agent Circuit ID and Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.

Options

`circuit-id`—(Optional) Agent Circuit ID suboption (suboption 1).

`remote-id`—(Optional) Agent Remote ID suboption (suboption 2).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

option-82 (DHCP Local Server Pool Matching)

IN THIS SECTION

- Syntax | 2116
- Hierarchy Level | 2116
- Description | 2117
- Required Privilege Level | 2117
- Release Information | 2117

Syntax

```
option-82;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Configure the extended DHCP local server to use the option 82 value in the DHCP client DHCP PDU together with the ip-address-first method to determine which address-assignment pool to use. You must configure the `ip-address-first` statement before configuring the `option-82` statement. The DHCP local server first determines which address-assignment pool to use based on the ip-address-first method. Then, the local server matches the option 82 value in the client PDU with the option 82 configuration in the address-assignment pool. This statement is supported for IPv4 address-assignment pools only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | **773**

Understanding Differences Between Legacy DHCP and Extended DHCP | **691**

Address-Assignment Pools Overview | **1139**

option-82 (Address-Assignment Pools)

IN THIS SECTION

- Syntax | **2118**
- Hierarchy Level | **2118**
- Description | **2118**
- Options | **2118**
- Required Privilege Level | **2119**

- Release Information | 2119

Syntax

```
option-82 {
    circuit-id value range named-range;
    remote-id value range named-range;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes option-match],
[edit access protocol-attributes attribute-set-name option-match]
```

Description

Specify the list of option 82 suboption match criteria used to select the named address range used for the client. The server matches the option 82 value in the user PDU to the specified option 82 match criteria and uses the named address range associated with the string.

Options

circuit-id Configure the address-assignment pool named-range to use for a particular option 82 Agent Circuit ID value.

- Values:
  - *value*—String for the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) in DHCP packets.
  - range *named-range*—Name of the address-assignment pool range to use.

remote-id Specify the address-assignment pool named range to use based on the particular option 82 Agent Remote ID value.

- Values:
  - range *named-range*—Name of the address-assignment pool range to use.
  - *value*—String for Agent Remote ID suboption (suboption 2) of the DHCP relay agent information option (option 82) in DHCP packets.

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 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

option-match

IN THIS SECTION

- Syntax | 2120
- Hierarchy Level | 2120
- Description | 2120
- Required Privilege Level | 2120
- Release Information | 2120

Syntax

```
option-match {
    option-82 {
        circuit-id value range named-range;
        remote-id value range named-range;
    }
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes],
[edit access protocol-attributes attribute-set-name]
```

Description

Specify a list of match criteria used to determine which named address range in the address-assignment pool to use. The extended DHCP local server matches this information to the match criteria specified in the client PDUs. For example, for option 82 match criteria, the server matches the option 82 value in the user PDU to the specified option 82 string and uses the named range associated with the string.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

option-number (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 2121
- Hierarchy Level | 2121
- Description | 2121
- Options | 2122
- Required Privilege Level | 2122
- Release Information | 2122

Syntax

```
option-number option-number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay group group-name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option],
[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 ...]
```

Description

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

Use the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level to configure the DHCPv6 relay agent support.

Options

option-number—The DHCP or DHCPv6 option in the incoming traffic.

NOTE: EX Series switches do not support the User Class Options.

- 15 (DHCPv6 only)—Use DHCPv6 option 15 (User Class Option) in packets
- 16 (DHCPv6 only)—(MX Series routers and EX Series switches only) Use DHCPv6 option 16 (Vendor Class Option) in packets
- 60 (DHCPv4 only)—(MX Series routers and EX Series switches only) Use DHCP option 60 (Vendor Class Identifier) in DHCP packets
- 77 (DHCPv4 only)—Use DHCP option 77 (User Class Identifier) in packets

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

options (Access Profile)

IN THIS SECTION

Syntax | 2123
Hierarchy Level | 2125

- Description | 2125
- Options | 2125
- Required Privilege Level | 2131
- Release Information | 2132

Syntax

```
options {
    accounting-session-id-format (decimal | description);
    calling-station-id-delimiter delimiter-character;
    calling-station-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    chap-challenge-in-request-authenticator;
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm  (direct | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    interface-description-format {
        exclude-adapter;
        exclude-channel;
        exclude-sub-interface;
    }
    ip-address-change-notify message;
    juniper-access-line-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
```

```
            port-width width;
            pw-width width;
            slot-width width;
            vci-width width;
            vpi-width width;
        }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
        agent-circuit-id;
        agent-remote-id;
        concatenated-vlan-tags  {
            fixed-size-inner-tag;
            fixed-size-outer-tag;
        }
        interface-description;
        interface-text-description;
        nas-identifier;
        order {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            postpend-vlan-tags;
        }
        postpend-vlan-tags;
    }
    nas-port-type {
        ethernet {
            port-type;
        }
    }
    override {
        calling-station-id remote-circuit-id;
        nas-ip-address tunnel-client-gateway-address;
        nas-port tunnel-client-nas-port;
        nas-port-type tunnel-client-nas-port-type;
    }
    remote-circuit-id-delimiter;
    remote-circuit-id-fallback;
    remote-circuit-id-format {
        agent-circuit-id;
```

```
        agent-remote-id;
    }
    revert-interval interval;
    service-activation {
        dynamic-profile (optional-at-login | required-at-login);
        extensible-service (optional-at-login | required-at-login);
    }
    vlan-nas-port-stacked-format;
}
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Configure the options used by RADIUS authentication and accounting servers.

Options

accounting-session-id-format (EX Series, MX Series only) Configure the format the router or switch uses to identify the accounting session. The default is `decimal`.

- Values:
  - decimal—Use the decimal format.
  - description—Use the generic format, in the form: `jnpr` *interface-specifier*:*subscriber-session-id*.

calling-station-id-delimiter (MX Series, T Series only) Starting in Junos OS Release 13.1, specify the character that the router uses as a separator between the concatenated values in the Calling-Station-ID (RADIUS IETF attribute 31) string. The router uses the delimiter when you configure more than one value in the `calling-station-id-format` statement. The default is the hash (#) character.

- Values:
  - *delimiter-character*—Character to use for the delimiter. You must enclose the delimiter character in quotation marks (“ ”).

chap-challenge-in-request-authenticator

(MX Series only) Starting in Junos OS Release 15.1, configure the `authd` process to insert the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets, if the challenge value is 16 bytes long. If you enable the `chap-challenge-in -request-authenticator` statement and the random challenge is not 16 bytes long, `authd` ignores the statement and uses the default behavior, which inserts the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets.

client-accounting-algorithm

(EX Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the access method the router uses to access RADIUS accounting servers. The default is the `direct` option.

- Values:
  - direct—Use the direct method.
  - round-robin—Use the round-robin method.

client-authentication-algorithm

(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the method that the authenticator uses to access RADIUS authentication servers when there are multiple servers configured. Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The router or switch, acting as the authenticator, waits for a response from the server before sending another request.

When there are multiple RADIUS server connections configured for a client, the authenticator attempts to reach the different servers in the order that they are configured. If there is no response from the first RADIUS server, the authenticator attempts to reach the next RADIUS server. This process repeats until the client is either granted access or there are no more configured servers.

If the `direct` method is configured, the authenticator always treats the first server in the list as the primary server. The authenticator moves on to the second server only if the attempt to reach the first server fails. If the `round-robin` method is configured, the server chosen first will be rotated based on which server was used last. The first server in the list is treated as a primary for the first authentication request, but for the second request, the second server configured is treated as primary, and so on. With this method, all of the configured servers receive roughly the same number of requests on average so that no single server has to handle all of the requests.

NOTE: The `round-robin` access method is not supported on EX Series switches.

- **Default:** The default is the `direct` option.
- Values:
  - direct—Use the direct access method. The authenticator contacts the first RADIUS server on the list for each request, the second server if the first one fails, and so on.
  - round-robin—Use the round-robin method. The authenticator contacts the first RADIUS server for the first request, the second server for the second request, and so on.

coa-dynamic-variable-validation

(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, specify that when a CoA operation includes a change to a client profile dynamic variable that cannot be applied (such as an update to a non-existent filter), the router does not apply any changes to client profile dynamic variables in the request, and responds with a NACK message.

- **Default:** If you do not configure this statement, the router does not apply any incorrect variable updates, but does make any other changes to the client profile dynamic variables, and responds with an ACK message.

ethernet-port-type-virtual

(EX Series, M Series, MX Series only) Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of `ethernet` in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of `virtual`.

> **NOTE**: This statement takes precedence over the `nas-port-type` statement if you include both statements in the same access profile.

access-loop-id-local

Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

ip-address-change-notify

(MX Series only) Starting in Junos OS Release 13.1, for on-demand address allocation for dual-stack PPP subscribers, specify that the BNG includes the IPv4-Release-Control VSA (26–164) in the Access-Request that is sent during on-demand IP address allocation and in the Interim-Accounting messages that are sent to report an address change. The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

Optionally, configure a message that is included in the VSA when it is sent to the RADIUS server.

- **Default:** This functionality is disabled by default.
- **Values:** ***message***—VSA message.
- **Range:** Up to 32 characters.

juniper-access-line-attributes

Configure AAA to add Juniper Networks access line VSAs to the RADIUS authentication and accounting request messages for subscribers. If the router has not received and processed the corresponding ANCP attributes from the access node, then AAA provides only the following in these RADIUS messages:

- Downstream-Calculated-QoS-Rate (IANA 4874, 26-141)—Default configured advisory transmit speed.
- Upstream-Calculated-QoS-Rate (IANA 4874, 26-142)—Default configured advisory receive speed.

NOTE: Starting in Junos OS Release 19.2R1, the `juniper-access-line-attributes` option replaces the `juniper-dsl-attributes` option. The difference between these options is that `juniper-dsl-attributes` supported only DSL TLVs received in the ANCP Port Status message. The `juniper-access-line-attributes` option supports PON TLVs in addition to DSL TLVs, and will be extensible to future access technologies.

For backward compatibility with existing scripts, the `juniper-dsl-attributes` option redirects to the new `juniper-access-line-attributes` option. We recommend that you use `juniper-access-line-attributes`.

NOTE: The `juniper-access-line-attributes` option is not backward compatible with Junos OS Release 19.1 or earlier releases. This means that if you have configured `juniper-access-line-attributes` option in Junos OS Release 19.2 or higher releases, you must perform the following steps to downgrade to Junos OS Release 19.1 or earlier releases:

1. Delete the `juniper-access-line-attributes` option from all access profiles that include it.
2. Perform the software downgrade.
3. Add the `juniper-dsl-attributes` option to the affected access profiles.

- **Default:** The Juniper Networks access line VSAs are not added to the RADIUS authentication and accounting request messages. However, the DSL Forum VSA —if available—is added to RADIUS messages by default.

nas-identifier

(EX Series, MX Series, SRX Series only) Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests. This statement was introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

- **Values:** *identifier-value*—String to use for authentication and accounting requests.
- **Range:** 1 through 64 characters.

nas-port-id-delimiter

(MX Series only) Starting in Junos OS Release 11.4, specify the character that the router uses as a separator between the concatenated values in the NAS-Port-ID string. The router uses the delimiter when you configure more than one value in the `nas-port-id-format` statement. The default is the hash (#) character. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

- **Values:** *delimiter-character*—Character used for the delimiter.

remote-circuit-id-delimiter

(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure a delimiter character for the remote circuit ID string when you use the `remote-circuit-id-format` statement to configure the string to use instead of the Calling-Station ID in L2TP Calling Number AVP 22. If more than one value is configured for the remote circuit ID format, the delimiter character is used as a separator between the concatenated values in the resulting remote circuit ID string. The default is the hash (#) character.

- **Values:** *delimiter*—Delimiter character to be used between components of the remote circuit ID string.

remote-circuit-id-fallback

(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the fallback value for the LAC to send in L2TP Calling Number AVP 22, either the configured Calling-Station-ID or the default underlying interface. Use of the fallback value is triggered when the components of the override string you configured with the `remote-circuit-id-format` statement—the ACI, the ARI, or both ACI and ARI—are not received by the LAC in the PPPoE Active Discovery Request (PADR) packet.

- Values:

- configured-calling-station-id—Send the configured Calling-Station-ID in the Calling Number AVP.
- default—Send the underlying interface value in the Calling Number AVP.

remote-circuit-id-format (MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the format of the string that overrides the Calling-Station-ID format in the Calling Number AVP 22 sent by the LAC to the LNS in the ICRQ packet when an L2TP session is being established. You can specify the ACI, the ARI, or both the ACI and ARI. This statement enables you to decouple the AVP 22 value from the RADIUS Calling-Station-ID attribute (31); the values for AVP 22 and the Calling-Station-ID attribute are the same when you use the `calling-station-id-format` statement to configure AVP 22.

> **NOTE**: You must configure the `override calling-circuit-id remote-circuit-id` statement for the remote circuit ID format to be used in the calling number AVP.

- Values:
  - agent-circuit-id—Specifies use of the ACI string that uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. For PPPoE traffic, the ACI string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.
  - agent-remote-id—Specifies use of the ARI string that identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The agent remote identifier (ARI) string is stored in the DSL Forum Agent-Remote-ID VSA [26-2] for PPPoE traffic.

service-activation (MX Series only) Starting in Junos OS Release 16.2, specify whether subscribers are allowed to log in even when service activation failures related to configuration errors occur during family activation request processing by authd for a newly authenticated subscriber. Configuration errors include missing or incorrect syntax, missing or incomplete references to dynamic profiles, and missing or incomplete variables.

NOTE: This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policies.

You can enable separate configurations for subscriber login services for two `service-activation` types: `dynamic-profile` and `extensible-service`. You configure the `dynamic-profile` type services in the dynamic profile at the `[edit dynamic-profiles]` hierarchy level; the profile is used to provide dynamic subscriber access and services for broadband applications. The `extensible-service` type is for business services configured in an operation script and provisioned by the Extensible Subscriber Services Manager daemon (essmd).

- Default:

  Default behavior depends on the service type:

  - For `extensible-service` services: `optional-at-login`.
  - For `dynamic-profile` services: `required-at-login`.

- Values:

  - optional-at-login—Service activation is optional. Failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Failure for any other reason causes network family activation to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.
  - required-at-login—Service activation is required. Failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.

vlan-nas-port-stacked-format (MX Series only) Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`juniper-dsl-attributes` introduced in Junos OS Release 11.4.

`nas-port-id-delimiter` introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

`calling-station-id-delimiter` introduced in Junos OS Release 13.1.

`ip-address-change-notify` introduced in Junos OS Release 13.1.

`coa-dynamic-variable-validation`, `client-authentication-algorithm`, and `client-accounting-algorithm` introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

`remote-circuit-id-delimiter`, `remote-circuit-id-fallback`, and `remote-circuit-id-format` introduced in Junos OS Release 13.3R1 on MX Series.

`chap-challenge-in-request-authenticator` introduced in Junos OS Release 15.1.

`nas-identifier` introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

`service-activation` introduced in Junos OS Release 16.2.

`juniper-access-line-attributes` introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

RADIUS Authentication and Accounting Basic Configuration | 550

order

IN THIS SECTION

Syntax | 2133

- Hierarchy Level | 2133
- Description | 2133
- Options | 2133
- Required Privilege Level | 2133
- Release Information | 2133

Syntax

```
order [ accounting-method ];
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Set the order in which the Junos OS tries different accounting methods for client activity. When a client logs in, the software tries the accounting methods in the specified order.

Options

accounting-method—One or more accounting methods. When a client logs in, the software tries the accounting methods in the following order, from first to last. The only valid value is `radius` for RADIUS accounting.

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 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration | 550

origin (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 2134
- Hierarchy Level | 2134
- Description | 2134
- Options | 2135
- Required Privilege Level | 2135
- Release Information | 2135

Syntax

```
origin realm realm-name host hostname
```

Hierarchy Level

```
[edit diameter]
```

Description

Specify the hostname and realm of the endpoint node that originates Diameter messages for the Diameter instance. The Diameter instance supplies these values to be conveyed by the Origin-Realm-AVP and Origin-Host-AVP for all messages sent by the Diameter instance.

NOTE: Both the host and realm are mandatory.

Options

host *hostname* Name of the message origin host that is supplied as the value of the Origin-Host AVP for all Diameter messages in the Diameter master instance.

realm *realm-name* Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for all Diameter messages in the Diameter master instance.

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring the Origin Attributes of the Diameter Instance | 1374

other-bytes

IN THIS SECTION

- Syntax | 2136
- Hierarchy Level | 2136
- Description | 2136
- Options | 2136
- Required Privilege Level | 2136
- Release Information | 2137

Syntax

```
other-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an access line of DSL type OTHER. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `other-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Configuring the ANCP Agent | **1253**

other-overhead-adjust

IN THIS SECTION

- Syntax | **2137**
- Hierarchy Level | **2138**
- Description | **2138**
- Options | **2138**
- Required Privilege Level | **2138**
- Release Information | **2138**

Syntax

```
other-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for an access line of DSL type OTHER received in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `other-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Configuring the ANCP Agent | **1253**

other-stateful-configuration (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | **2139**
- Hierarchy Level | **2139**
- Description | **2140**
- Default | **2140**
- Required Privilege Level | **2140**
- Release Information | **2140**

Syntax

```
(other-stateful-configuration | no-other-stateful-configuration);
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable autoconfiguration of other nonaddress-related information:

- `no-other-stateful-configuration`—Disable autoconfiguration of other nonaddress-related information.
- `other-stateful-configuration`—Enable autoconfiguration of other nonaddress-related information.

Default

The configured object is disabled unless explicitly enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

overhead-accounting (ANCP)

IN THIS SECTION

- Syntax | 2141
- Hierarchy Level | 2141
- Description | 2141
- Required Privilege Level | 2141
- Release Information | 2141

Syntax

```
overhead-accounting;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-name]
```

Description

Prevent ANCP from performing an adjustment on the actual downstream data rate that ANCP receives from the DSLAM for the difference between the customer premise equipment (CPE) protocol overhead and the B-RAS protocol overhead. You include this statement when you want CoS to perform the adjustment on the data rate from the DSLAM according to the overhead accounting configuration in a CoS traffic control profile.

When this statement is not configured (the default condition), ANCP makes the traffic rate adjustment according to the configuration of the `qos-adjust-line-type` statements and reports that rate to CoS. CoS then applies (if configured) the adjustment set by the `overhead-accounting` statement in the CoS traffic profile.

NOTE: ANCP reports a traffic rate to CoS only if the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level has been configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

override-chap-password

IN THIS SECTION

- Syntax | 2142
- Hierarchy Level | 2142
- Release Information | 2142
- Description | 2142
- Required Privilege Level | 2143

Syntax

```
override-chap-password password;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

Configure a CHAP password to override the existing password for authenticating any subscriber associated with the domain map. The override CHAP password replaces the CHAP challenge response that the PPP client sends to the server for authentication.

The CHAP challenge response is normally sent in the RADIUS Access-Request message as the CHAP-Password attribute (attribute 3). When you configure an override CHAP password, it is sent in the Access-Request as the User-Password attribute (attribute 2), and the CHAP-Password attribute is not included in the message.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement in the configuration.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning | 671

override-password (Domain Map)

IN THIS SECTION

- Syntax | 2143
- Hierarchy Level | 2143
- Description | 2144
- Options | 2144
- Required Privilege Level | 2144
- Release Information | 2144

Syntax

```
override-password password
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Specify a password to be used as the authenticating password for all outgoing authentication requests for subscribers who match the domain map.

NOTE: This override works only when PAP is the authentication method. It does not work for CHAP.

Options

password	Name of the password.

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 14.1.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning | 671

overrides (DHCP Local Server)

IN THIS SECTION

- Syntax | 2145
- Hierarchy Level | 2146
- Description | 2146

Required Privilege Level | 2147

Release Information | 2147

Syntax

```
asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
```

```
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delegated-pool;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Override the default configuration settings for the extended DHCP local server. Specifying the `overrides` statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.

- To override global 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 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 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.
- Use the `[edit system services dhcp-local-server dhcpv6]` hierarchy level to override DHCPv6 configuration options.

NOTE: By default, `jdhcp` does not process DHCPINFORM message. Only after you enable the overrides command using the `set system services dhcp-local-server overrides process-inform` statement, `jdhcp` starts processing the DHCPINFORM message.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

The `interface-client-limit` statement is not supported in the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

The `asymmetric-prefix-lease-time`, `delegated-pool`, `multi-address-embedded-option-response`, and `rapid-commit` statements are supported in the `[edit system services dhcp-local-server dhcpv6 ...]` hierarchy level only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the `allow-no-end` option introduced in Junos OS Release 14.1X53-D15 for EX Series switches.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP | 691

Overriding the Default DHCP Local Server Configuration Settings | 707

Configuring a DHCP Server on Switches

overrides (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2148
- Hierarchy Level | 2149
- Description | 2149
- Required Privilege Level | 2150
- Release Information | 2150

Syntax

```
overrides {
    allow-no-end-option;
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
```

```
    send-release-on-delete;
    trust-option-82;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Override the default configuration settings for the extended DHCP relay agent. Specifying the `overrides` statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

The following statements are supported at both the `[edit ... dhcp-relay]` and `[edit ... dhcpv6]` hierarchy levels.

- `allow-snooped-clients`
- `asymmetric-lease-time`
- `delete-binding-on-renegotiation`
- `dual-stack`
- `interface-client-limit`
- `no-allow-snooped-clients`
- `no-bind-on-request`
- `relay-source`

- send-release-on-delete

The following statements are supported at the [edit ... dhcpv6] hierarchy levels only.

- asymmetric-prefix-lease-time

All other statements are supported at the [edit ... dhcp-relay] hierarchy levels only.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Overriding the Default DHCP Relay Configuration Settings | 709

overrides (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 2151
- Hierarchy Level | 2152
- Description | 2152
- Options | 2152
- Required Privilege Level | 2153

- Release Information | 2153

Syntax

```
overrides {
    event {
        catastrophic-failure {
            reboot (master | standby);
        }
    }
    interfaces {
        family (inet | inet6) {
            layer2-liveness-detection;
            ipoe-dynamic-arp-enable;
            receive-gratuitous-arp;
        }
    }
    no-unsolicited-ra;
    ra-initial-interval-max seconds;
    ra-initial-interval-min seconds;
    shmlog {
        disable;
        file filename <files maximum-no-files> <size maximum-file-size>;
        filtering enable;
        log-name {
            all;
            logname {
                <brief | detail | extensive | none | terse>;
                <file-logging |no-file-logging>;
            }
            }
        log-type (debug | info | notice);
    |
}
```

Hierarchy Level

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

Description

Override the default configuration settings for the Junos OS enhanced subscriber management software for subscriber management.

Options

ra-initial-interval-max *seconds*

Specify the high end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the `ra-initial-interval-min` option.

- **Range:** 1 through 16

ra-initial-interval-min *seconds*

Specify the low end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the `ra-initial-interval-max` option.

> **BEST PRACTICE**: Always configure the value of `ra-initial-interval-min` to be less than or equal to the value of `ra-initial-interval-max`. If you configure the values to be the same, the initial router advertisement intervals are constant and not randomized.

- **Range:** 1 through 16

ipoe-dynamic-arp-enable

Enable dynamic ARP to resolve the MAC address for IPv4 framed host (32-bit) routes. By default the framed route is permanently associated with the source MAC address received in the packet that triggered creation of the dynamic VLAN.

receive-gratuitous-arp

Enable the router to compare the source MAC address received in a gratuitous ARP request or reply packet with the value in the ARP cache. The router updates the cache with the received MAC address when it determines this address is different from the cache entry.

This situation occurs when an IPv4 address is moved to a different device. The device broadcasts a gratuitous ARP reply packet with its MAC address as the source MAC

address. When the receive-gratuitous-arp option is configured, the router compares the MAC addresses and updates the cache to associate the IPv4 address with the new MAC address.

If the receive-gratuitous-arp option is not configured, the router does not accept the gratuitous ARP request or reply packet and cannot quickly learn about the new address. Instead, the original dynamic ARP entry in the cache eventually times out. Before deleting the entry, the router sends an ARP request for the target IP address. The client responds with the new MAC address. This delay in learning about the new address means there is a period during which the MAC address in the ARP cache does not match the address in the new device's NIC.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

ra-initial-interval-max and ra-initial-interval-min options added in Junos OS Release 18.2R1 on MX Series routers.

ipoe-dynamic-arp-enable and receive-gratuitous-arp options added in Junos OS Release 18.4R1 on MX Series routers.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors | **942**

parse-direction (Domain Map)

IN THIS SECTION

- Syntax | 2154
- Hierarchy Level | 2154
- Description | 2154
- Default | 2154
- Options | 2154
- Required Privilege Level | 2155
- Release Information | 2155

Syntax

```
parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the direction in which the router searches for the domain name in a username.

Default

`left-to-right`

Options

`left-to-right`—The router searches starting at the left-most character. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.

`right-to-left`—The router searches starting at the right-most character. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.

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 the Parsing Direction for Domain and Realm Names | 670

Configuring Domain and Realm Name Usage for Domain Maps | 667

parse-order (Domain Map)

IN THIS SECTION

- Syntax | 2155
- Hierarchy Level | 2156
- Description | 2156
- Default | 2156
- Options | 2156
- Required Privilege Level | 2156
- Release Information | 2156

Syntax

```
parse-order (domain-first | realm-first);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the order in which the router searches for a domain name, either the domain first, or the realm first.

Default

`domain-first`

Options

`domain-first`—The router searches for a domain name first. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.. If no domain is found, then the router searches for a realm. If the router does not find either a domain or realm, then there is no domain.

`realm-first`—The router searches for a realm name first. When the router reaches a realm delimiter, it uses anything to the left of the delimiter as the domain name. If no realm is found, then the router searches for a domain. If the router does not find either a domain or realm, then there is no domain.

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 14.1.

RELATED DOCUMENTATION

Specifying the Parsing Direction for Domain and Realm Names | 670

Configuring Domain and Realm Name Usage for Domain Maps | 667

partition

IN THIS SECTION

- Syntax | 2157
- Hierarchy Level | 2157
- Description | 2157
- Options | 2157
- Required Privilege Level | 2158
- Release Information | 2158

Syntax

```
partition partition-name  {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm;
}
```

Hierarchy Level

```
[edit jsrc]
```

Description

Configure a JSRC partition.

Options

partition-name—Name of the JSRC partition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview | 1476

Configuring the JSRC Partition | 1477

partition (Gx-Plus)

IN THIS SECTION

- Syntax | 2158
- Hierarchy Level | 2159
- Description | 2159
- Options | 2159
- Required Privilege Level | 2159
- Release Information | 2159

Syntax

```
partition partition-name  {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm;
}
```

Hierarchy Level

```
[edit access gx-plus]
```

Description

Configure a Gx-Plus partition.

Options

partition-name—Name of the Gx-Plus partition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus | 1404

Configuring the Gx-Plus Partition | 1405

partition (NASREQ Diameter Application)

IN THIS SECTION

Syntax | 2160

- Hierarchy Level | 2160
- Description | 2160
- Options | 2160
- Required Privilege Level | 2161
- Release Information | 2161

Syntax

```
partition partition-name {
    destination-host hostname;
    destination-realm realm-name;
    diameter-instance master;
}
```

Hierarchy Level

```
[edit access nasreq]
```

Description

Define the NASREQ partition by specifying the destination host and realm where the application resides. Only the`master` routing instance is supported for the application partition.

Options

diameter-instance master	Specifies that the master Diameter instance is used. `master` is the only supported value.
hostname	Name of the host where the NASREQ server application resides; generally, this value is not set unless needed.
partition-name	Name of the NASREQ partition.
realm-name	Name of the realm where the NASREQ server application resides.

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 16.1.

RELATED DOCUMENTATION

Diameter Network Access Server Application (NASREQ) | 1463

Configuring the Diameter Network Access Server Application (NASREQ) | 1465

partition (OCS)

IN THIS SECTION

- Syntax | 2161
- Hierarchy Level | 2162
- Description | 2162
- Options | 2162
- Required Privilege Level | 2163
- Release Information | 2163

Syntax

```
partition partition-name  {
    called-station-id station-name;
    charging-id number;
    destination-host ocs-hostname;
    destination-realm ocs-realm-name;
    diameter-instance;
```

```
    draining;
    draining-response-timeout seconds;
    final-response-timeout seconds;
    force-continue;
    ggsn-address address;
    ggsn-mcc-mncggsn-mcc-mnc;
    max-outstanding-requests number;
    send-origin-state-id number;
    user-name-include  {
        base-interface-name;
        delimiter delimiter-character;
        domain-name domain-name;
        interface-name;
        mac-address;
        nas-port-id;
        origin-host;
        origin-realm;
        user-name;
        user-prefix prefix;
    }
}
```

Hierarchy Level

```
[edit access ocs]
```

Description

Configure an Online Charging System (OCS) partition (a specific logical system:routing instance context) and its parameters to define user ID information, destination host and realm names, and number of outstanding requests to the OCS, and control subscriber traffic flow before the first interrogation with the OCS. The OCS that interacts with the Policy and Charging Enforcement Function (PCEF). Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

Options

partition-name—Name of the OCS partition. You can define only one partition for the OCS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

partition (PCRF)

IN THIS SECTION

- Syntax | **2164**
- Hierarchy Level | **2165**
- Description | **2165**
- Options | **2165**
- Required Privilege Level | **2166**
- Release Information | **2166**

Syntax

```
partition partition-name  {
    destination-host pcrf-hostname;
    destination-realm pcrf-realm-name;
    diameter-instance;
    draining;
    draining-response-timeout seconds;
    ip-can-type number;
    local-decision {
        deny;
        grant;
        reinit-on-failure;
        reinit-on-rar;
        reinit-timeout seconds;
        timeout seconds;
    }
    logout-response-timeout seconds;
    max-outstanding-requests number;
    report-local-rule;
    report-resource-allocation;
    report-successful-resource-allocation;
    send-dyn-subscription-indicator;
    send-network-family-indicator;
    send-origin-state-id;
    subscription-id-data-include  {
        base-interface-name;
        delimiter delimiter-character;
        domain-name name;
        interface-name;
        mac-address;
        nas-port-id;
        origin-host;
        origin-realm;
        user-name;
        user-prefix prefix;
        vlan-tags;
    }
    subscription-id-type number;
    update-response-timeout seconds;
```

```
    use-session-stamp;
}
```

Hierarchy Level

```
[edit access pcrf]
```

Description

Configure a Policy and Charging Rules Function (PCRF) partition (a specific logical system:routing instance context) and its parameters to define subscription ID information, destination host and realm names, and reporting for rules and resources. 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.

Options

partition-name Name of the PCRF partition. You can define only one partition for the PCRF.

use-session-stamp (Optional) Adds a 32-bit string to the end of the PCRF session ID for new subscribers. Existing subscribers are not affected when you configure this option. The string, called the session stamp or the session timestamp, consists of the UTC time when the router creates the CCR-GX-I. Appending the timestamp is considered to be an extended format for the session ID. You can use the extended-format session ID when you need an eternally unique ID or when you configure local reinitialization.

> **NOTE**: You must configure the `use-session-stamp` option when you configure local reinitialization with the `reinit-on-failure` or `reinit-on-rar` options with the "local-decision" on page 2008 statement.

> **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.

Click a linked statement in the Syntax section for more information about that statement.

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 16.2.

`use-session-stamp` option added in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | **1455**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**
- Understanding Upstream and Downstream Messages for the PCRF | **1444**

partition (s6a)

IN THIS SECTION

- Syntax | **2167**
- Hierarchy Level | **2167**
- Description | **2167**
- Options | **2167**
- Required Privilege Level | **2168**
- Release Information | **2168**

Syntax

```
partition name {
    destination-host destination-host;
    destination-realm destination-realm;
    diameter-instance diameter-instance;
    max-outstanding-requests max-outstanding-requests;
    response-timeout seconds;
}
```

Hierarchy Level

```
[edit access s6a]
```

Description

Define the s6a partition by specifying the destination host and realm where the application resides and specify the endpoint origin, the remote peers, and the network elements that associate routes with peers to control diameter forwarding of S6a messages. Only the `master` routing instance is supported for the application partition.

Options

name — Name of the S6a partition.

destination-host — Name of the host where the S6A server application resides; generally, this value is not set unless needed.

destination-realm — Name of the realm where the S6a server application resides.

diameter-instance — Specifies that the master Diameter instance is used. `master` is the only supported value.

max-outstanding-requests — Configure the maximum number of outstanding requests for the S6a application.

- **Default:** 40
- **Range:** 2 through 1024

response-timeout Configure the amount of time in seconds that the Mobility Management Entity (MME) waits to receive a response from the Home Subscriber Server (HSS).

- **Default:** 15
- **Range:** 1 through 30

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- Example: Configure S6a Application | **1379**
- authentication-order | **1640**
- Diameter Base Protocol Overview | **1339**

password (Static Subscribers)

IN THIS SECTION

- Syntax | **2169**
- Hierarchy Level | **2169**
- Description | **2169**
- Options | **2169**
- Required Privilege Level | **2169**
- Release Information | **2169**

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers authentication],
[edit logical-systems logical-system-name system services static-subscribers authentication],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication],
[edit routing-instances routing-instances-name system services static-subscribers authentication],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],  authentication],
[edit system services static-subscribers authentication],
[edit system services static-subscribers group group-name authentication]
```

Description

Specify the password that is sent to AAA for user login for all static subscribers on interfaces configured at the [edit system services static-subscribers interface] hierarchy level, or for the subscribers in a specified group. The group version of the statement takes precedence over the global version.

Options

password-string—String that defines the password.

Required Privilege Level

system-level—To view this statement in the configuration.

system-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 | 1487

Configuring the Static Subscriber Global Authentication Password | 1492

Configuring the Static Subscriber Group Authentication Password | 1497

password (DHCP Local Server)

IN THIS SECTION

- Syntax | 2170
- Hierarchy Level | 2170
- Description | 2171
- Options | 2171
- Required Privilege Level | 2171
- Release Information | 2172

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
```

```
authentication],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication],
[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name authentication],
[edit system services dhcp-local-server group group-name authentication]
```

Description

Configure the password that is sent to the external AAA authentication server for subscriber authentication or DHCP client authentication.

Options

password-string—Authentication password.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

password (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2172
- Hierarchy Level | 2172
- Description | 2173
- Options | 2173
- Required Privilege Level | 2173
- Release Information | 2173

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication],
[edit forwarding-options dhcp-relay dhcpv6 authentication],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication],
[edit forwarding-options dhcp-relay group group-name authentication],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication]
```

Description

Configure the password that is sent to the external AAA authentication server for subscriber authentication or client authentication. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

password-string—Authentication password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the [edit ... dual-stack-group *dual-stack-group-name* authentication] hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Example-Configuring DHCP with External Authentication Server | 837

pcrf (Diameter Applications)

IN THIS SECTION

- Syntax | 2174
- Hierarchy Level | 2175
- Description | 2175
- Required Privilege Level | 2176
- Release Information | 2176

Syntax

```
pcrf {
    global {
        rule-param avp-code;
    }
    partition partition-name  {
        destination-host pcrf-hostname;
        destination-realm pcrf-realm-name;
        diameter-instance;
        draining;
        draining-response-timeout seconds;
        ip-can-type number;
        local-decision {
            deny;
            grant;
```

```
            reinit-on-failure;
            reinit-on-rar;
            reinit-timeout seconds;
            timeout seconds;
        }
        logout-response-timeout seconds;
        max-outstanding-requests number;
        report-local-rule;
        report-resource-allocation;
        report-successful-resource-allocation;
        send-dyn-subscription-indicator;
        send-network-family-indicator;
        send-origin-state-id;
        subscription-id-data-include  {
            base-interface-name;
            delimiter delimiter-character;
            domain-name name;
            interface-name;
            mac-address;
            nas-port-id;
            origin-host;
            origin-realm;
            user-name;
            user-prefix prefix;
            vlan-tags;
        }
        subscription-id-type number;
        update-response-timeout seconds;
        use-session-stamp;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Policy and Charging Rules Function (PCRF) global attributes, rules, parameters, and partition to authorize and provision subscribers. PCRF is 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.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

peer (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 2177
- Hierarchy Level | 2177
- Description | 2177
- Options | 2177
- Required Privilege Level | 2177
- Release Information | 2177

Syntax

```
peer  peer-name {
    address ip-address;
    connect-actively {
        port port-number;
        transport transport-name;
    }
    logical-system logical-system-name  <routing-instance routing-instance-name>;
    peer-origin realm realm-name host hostname;
    routing-instance routing-instance-name;
    send-origin-state-id;
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Configure peer-specific Origin-Realm and Origin-Name AVP values to include in the CER, DWR, DWA, DPR, and DPA messages of this peer.

Options

peer-name—Name of the peer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

peer (Diameter Network Element)

IN THIS SECTION

- Syntax | 2178
- Hierarchy Level | 2178
- Description | 2178
- Options | 2179
- Required Privilege Level | 2179
- Release Information | 2179

Syntax

```
peer peer-name {
    priority priority-value;
}
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Assigns a peer to a Diameter network element and sets a priority value for that peer within the DNE.

Options

peer-name—Name of the peer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

peer-ip-address-optional

IN THIS SECTION

- Syntax | 2180
- Hierarchy Level | 2180
- Description | 2180
- Required Privilege Level | 2180
- Release Information | 2180

Syntax

```
peer-ip-address-optional;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit interfaces interface-name pp0 unit unit-number ppp-options],
[edit access group-profile profile-name ppp ppp-options]
```

Description

Enable the Internet Protocol Control Protocol (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.

If the client’s provisioned IP address on the customer premises equipment (CPE):

- Matches the Framed-Route RADIUS attribute, then you must configure the access route at the `[edit dynamic-profiles routing-options]` hierarchy level.
- Matches the Framed-IP-Address RADIUS attribute, then no access route configuration is required.

NOTE: You must assign an IP address by configuring the Framed-IP-Address RADIUS attribute or the Framed-Pool RADIUS attribute, or by allocating an IP address from the local address pool without a RADIUS-specified pool name, with an *optional* Framed-Route RADIUS attribute returned from the RADIUS Server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

IPCP Negotiation with Optional Peer IP Address | 1120

Dynamic Profiles Overview

peer-origin (Diameter Peer)

IN THIS SECTION

- Syntax | 2181
- Hierarchy Level | 2181
- Description | 2181
- Options | 2182
- Required Privilege Level | 2182
- Release Information | 2182

Syntax

```
peer-origin realm realm-name host hostname;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify values of Origin-Realm-AVP and Origin-Host-AVP used in messages sent for the specified peer by the Diameter instance.

NOTE: Both the host and realm are mandatory for the peer origin.

Options

host *hostname*	Name of the message origin host that is supplied as the value of the Origin-Host AVP for Diameter messages to the peer.
realm *realm-name*	Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for Diameter messages to the peer.

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 16.2.

RELATED DOCUMENTATION

Configuring Diameter Peers | 1374

Configuring Diameter | 1373

pon (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 2183
- Hierarchy Level | 2184
- Description | 2184
- Options | 2185
- Required Privilege Level | 2186
- Release Information | 2187

Syntax

```
pon {
    gpon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    other {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    twdm-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    type tlv-value {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    wdm-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    xg-pon1 {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    xgs-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Configure adjustments to actual PON data rates as follows:

- Multiply the actual data rate by a percentage for only downstream rates or for both upstream an downstream rates
- Adjust the encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes

The actual (unadjusted) downstream and upstream data rates, PON line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

If the PON-Access-Type TLV (0x92) is not received in either the ANCP Port Status message or PPPoE-IA tags, the default adjustment leaves the rates and bytes unchanged.

Adjustments are applied to all subscribers using access lines of the specific PON line type. Depending on the value, it may be reported to AAA, CoS, or both:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate. A change in value triggers an immediate Interim-Accounting message to the RADIUS server if you have configured the "ancp-speed-change-immediate-update" on page 1621 statement.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following AVPs and messages:

- AVP 24, Tx Connect Speed (ICCN message)—For the initial total rate adjustment to ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94).

- AVP 38, Rx Connect Speed (ICCN message)—For the initial total rate adjustments to ONT/ONU-Peak-Data-Rate-Upstream TLV (0x95).
- AVP 97, Connect Speed Update (CSUN message)—For subsequent changes to the initial rates reported in AVP 24 and AVP 38.

Options

gpon Sets attributes for GPON access lines.

other Sets attributes for access lines of type OTHER. This might used, for example, when an OLT uses a new PON type that is not yet formally supported by the Broadband Forum.

overhead-adjust *percentage* Adjusts the actual downstream rates for all subscribers on an access line or PON tree of the specified types by multiplying the rate by the specified percentage. This adjustment accounts for the Layer 1 overhead for the PON type. The rates are adjusted as follows:

- For the subscriber access line, the adjustment is made to the downstream rate reported in the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94). The adjusted rate is applied to the logical interface or the (child) interface set for the subscriber access line, as determined by the CoS adjustment control profile. This adjustment applies for FTTH connections.
- For the PON tree, the adjustment is made to the downstream rate reported in the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98). The adjusted rate is applied to the PON tree (parent) interface set, as determined by the CoS adjustment control profile. This adjustment applies for both FTTH and FTTB connections. For FTTB, this is the rate on the OLT to DPU-P access line.
- **Range:** 80 through 100 percent
- **Default:** 100 percent

overhead-bytes *bytes* Adjusts the actual downstream cell overhead for all subscribers on the specified access line or PON tree by adding or subtracting the specified number of bytes. The adjustment accounts for the traffic encapsulation overhead. It shapes the logical interface or the (child) interface set for the subscriber access line or the PON tree (parent) interface set, as determined by the CoS adjustment control profile. The overhead is adjusted as follows:

- For the subscriber access line, the adjustment is made to the overhead reported in the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94). The adjusted overhead is

applied to the logical interface or the (child) interface set for the subscriber access line, as determined by the CoS adjustment control profile. This adjustment applies for FTTH connections.

- For the PON tree, the adjustment is made to the overhead reported in the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98). The adjusted overhead is applied to the PON tree (parent) interface set, as determined by the CoS adjustment control profile. This adjustment applies for both FTTH and FTTB connections. For FTTB, this is the overhead on the OLT to DPU-P access line.
- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

twdm-pon Sets attributes for TWDM-PON access lines.

total-adjust *percentage* Adjusts the downstream and upstream data rates for all subscribers on an access line of the specified type by multiplying the rate by the specified percentage. This adjustment accounts for the total Layer 1 and encapsulation overhead for the PON type. The adjustment is made to the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94) and the ONT/ONU-Maximum-Data-Rate-Upstream TLV (0x95). The adjustments apply to FTTH connections.

- **Range:** 1 through 100 percent
- **Default:** 100 percent

type *tlv-value* Sets attributes for access lines by specifying the unsigned integer value of the PON-Access-Type TLV (0x92) to reference the PON type. This option enables the `pon` statement to be used for PON access line types that might be introduced in the future.

- **Range:** 6 through 4294967295

wdm-pon Sets attributes for WDM-PON access lines.

xg-pon1 Sets attributes for XG-PON access lines.

xgs-pon Sets attributes for XGS-PON access lines.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates | 1308
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306
- Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299
- Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293
- Configuring the ANCP Agent | 1253

pool (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 2187
- Hierarchy Level | 2188
- Description | 2188
- Options | 2188
- Required Privilege Level | 2189
- Release Information | 2189

Syntax

```
pool pool-name {
    active-drain;
    family family {
        dhcp-attributes {
            [ protocol-specific attributes ]
        }
```

```
        excluded-address ip-address;
        excluded-range name low minimum-value high maximum-value;
        host hostname {
            hardware-address mac-address;
            ip-address ip-address;
        }
        network ip-prefixprefix-length>;
        prefix ipv6-prefix;
        range range-name {
            high upper-limit;
            low lower-limit;
            prefix-length prefix-length;
        }
    }
    hold-down;
    link pool-name;
}
```

Hierarchy Level

```
[edit access address-assignment]
[edit routing-instances routing-instances-name access address-assignment]
```

Description

Configure the name of an address-assignment pool.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

pool-name	Name assigned to the address-assignment pool.
active-drain	Configure the DHCP local server to stop allocating addresses from this pool. When this is configured, the DHCP local server gracefully shifts clients from this address pool to an alternative pool for which active drain is not configured. When existing clients with an

address from this pool submit a DHCPv4 request or DHCPv6 renew, they receive a NAK, forcing them to renegotiate. The server responds with a DHCPv4 offer or DHCPv6 advertise message with an address from a different pool.

family — Configure the protocol family for the address-assignment pool.

The options for this statement are explained separately. Click the linked statement for details.

hold-down — Configure an address-assignment pool that is currently in use to be unavailable for further address allocation. When a pool is in the hold-down state, the pool is no longer used to allocate IP addresses for subscribers. Current subscribers who previously obtained an address from the pool are not affected; they can continue to renew their leases. As each of these users disconnects, their address is not reallocated. The pool becomes inactive when all subscribers have disconnected and their addresses are returned to the pool.

link — Designate a secondary address-assignment pool that is linked to the pool being configured. When the pool being configured has no addresses available for allocation, the secondary pool can be searched for a free address. You can configure a chain of linked pools, but you cannot directly link more than one pool to or from any other pool. Each linked pool in the chain serves as a backup pool for the pool immediately before it in the chain.

- **Values:** *pool-name*—Name assigned to the secondary address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 9.0.

Support at the `[edit routing-instances routing-instances-name access address-assignment]` hierarchy level at tenant system level introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

DHCP Client Attribute and Address Assignment | 766

pool (DHCP Local Server Overrides)

IN THIS SECTION

- Syntax | 2190
- Hierarchy Level | 2190
- Description | 2191
- Options | 2192
- Required Privilege Level | 2192
- Release Information | 2192

Syntax

```
pool pool-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name overrides process-
inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name overrides process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server overrides process-
```

inform],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 overrides process-inform],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* overrides process-inform],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides process-inform],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* overrides process-inform],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* overrides process-inform],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides process-inform],
[edit system services dhcp-local-server overrides process-inform],
[edit system services dhcp-local-server dhcpv6 overrides process-inform],
[edit system services dhcp-local-server dhcpv6 group *group-name* overrides process-inform],
[edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides process-inform],
[edit system services dhcp-local-server group *group-name* overrides process-inform],
[edit system services dhcp-local-server group *group-name* interface *interface-name* overrides process-inform]

Description

Configure DHCP or DHCPv6 local server to reply to DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) with information taken from the specified pool without interacting with AAA.

Options

pool-name Name of the address pool, which must be configured within `family inet` for DHCP local server and within `family inet6` for DHCPv6 local server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Enabling Processing of Client Information Requests | 777

Overriding the Default DHCP Local Server Configuration Settings | 707

pool-match-order

IN THIS SECTION

- Syntax | 2193
- Hierarchy Level | 2193
- Description | 2193
- Default | 2193
- Required Privilege Level | 2193
- Release Information | 2193

Syntax

```
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

DHCP local server uses the `ip-address-first` method to determine which address pool to use.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use | 773

Understanding Differences Between Legacy DHCP and Extended DHCP | 691

port (Diameter Peer)

IN THIS SECTION

- Syntax | 2194
- Hierarchy Level | 2194
- Description | 2194
- Options | 2194
- Required Privilege Level | 2195
- Release Information | 2195

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit diameter peer peer-name connect-actively]
```

Description

Specify the destination TCP port used by the active connection to peer.

Options

port-number—Number of the TCP port.

- **Default:** 3868

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

pre-ietf-mode

IN THIS SECTION

- Syntax | 2195
- Hierarchy Level | 2196
- Description | 2196
- Required Privilege Level | 2196
- Release Information | 2196

Syntax

```
pre-ietf-mode
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Configure the ANCP agent to run in a mode that is backward compatible with Internet draft draft-wadhwa-gsmp-l2control-configuration-00.txt, *GSMP extensions for layer2 control (L2C)* for all neighbors or for a specific neighbor.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Configuring the ANCP Agent for Backward Compatibility | 1258

Configuring ANCP Neighbors | 1255

preauthentication-order (Access Profile)

IN THIS SECTION

- Syntax | 2197
- Hierarchy Level | 2197
- Description | 2197

Options | 2197
Required Privilege Level | 2197
Release Information | 2198

Syntax

```
preauthentication-order [ preauthentication-method ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which the Junos OS uses preauthentication methods for the LLID service when multiple methods are configured. Junos OS supports only the `radius` method.

NOTE: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the `exclude` statement.

Options

preauthentication-method

- `radius`—Verify the client using RADIUS.

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 13.3.

RELATED DOCUMENTATION

RADIUS Logical Line Identifier (LLID) Overview | 544

Configuring Logical Line Identification (LLID) Preauthentication | 546

preferred-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2198
- Hierarchy Level | 2198
- Description | 2199
- Options | 2199
- Required Privilege Level | 2199
- Release Information | 2199

Syntax

```
preferred-lifetime seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-
nameprefix prefix]
```

Description

Specify how long the prefix generated by stateless autoconfiguration remains preferred.

Options

seconds—Preferred lifetime, in seconds. If you set the preferred lifetime to `0xffffffff`, the lifetime is infinite. The preferred lifetime is never greater than the valid lifetime.

- **Default:** 604,800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

prefix (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2200
- Hierarchy Level | 2200
- Description | 2200
- Options | 2200
- Required Privilege Level | 2201
- Release Information | 2201

Syntax

```
prefix prefix;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)]
```

Description

Add a prefix to the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or to the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) information in DHCP packets that DHCP relay agent sends to a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name.

Options

prefix—Any of the following:

- `host-name`—Prepend the hostname of the router configured with the `host-name` statement at the `[edit system]` hierarchy level to the DHCP option information.
- `logical-system-name`—Prepend the name of the logical system to the option information.
- `routing-instance-name`—Prepend the name of the routing instance to the option information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... relay-agent-remote-id]` and `[edit ... remote-id]` hierarchy levels introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Using DHCP Relay Agent Option 82 Information | 751
- Using DHCP Relay Agent Option 82 Information | 751
- Configuring DHCPv6 Relay Agent Options | 917

prefix (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 2202
- Hierarchy Level | 2202
- Description | 2202
- Options | 2202
- Required Privilege Level | 2202
- Release Information | 2202

Syntax

```
prefix ipv6-prefix;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
```

Description

Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.

Options

ipv6-prefix—The IPv6 prefix.

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 10.0.

RELATED DOCUMENTATION

Address-Assignment Pools Overview | **1139**

Address-Assignment Pool Configuration Overview | **1146**

prefix (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2203
- Hierarchy Level | 2203
- Description | 2203
- Options | 2203
- Required Privilege Level | 2204
- Release Information | 2204

Syntax

```
prefix prefix {
    (autonomous | no-autonomous);
    (on-link | no-on-link);
    preferred-lifetime seconds;
    valid-lifetime seconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Configure the prefix name in router advertisement messages.

Options

prefix—Prefix name. For dynamic configuration, specify the *$junos-ipv6-ndra-prefix* dynamic variable.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

priority (Diameter Peer)

IN THIS SECTION

- Syntax | 2204
- Hierarchy Level | 2205
- Description | 2205
- Options | 2205
- Required Privilege Level | 2205
- Release Information | 2205

Syntax

```
priority priority-value;
```

Hierarchy Level

```
[edit diameter network-element element-name peer peer-name]
```

Description

Set the priority for a peer within a Diameter network element. A peer with a higher number has a higher priority.

Options

priority-value—Priority for the peer within the network element.

- **Range:** 1 through 65535

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

profile (Access)

IN THIS SECTION

Syntax | 2206

- Hierarchy Level | 2212
- Description | 2212
- Options | 2212
- Required Privilege Level | 2212
- Release Information | 2212

Syntax

```
profile profile-name {
    accounting {
        address-change-immediate-update
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        ancp-speed-change-immediate-update;
        coa-immediate-update;
        coa-no-override service-class-attribute;
        duplication;
        duplication-filter;
        duplication-vrf {
            access-profile-name profile-name;
            vrf-name vrf-name;
        }
        immediate-update;
        order [ accounting-method ];
        send-acct-status-on-config-change;
        statistics (time | volume-time);
        update-interval minutes;
        wait-for-acct-on-ack;
    }
    accounting-order (radius | [accounting-order-data-list]);
    authentication-order [ authentication-methods ];
    client client-name {
        chap-secret chap-secret;
        group-profile profile-name;
        ike {
            allowed-proxy-pair {
                remote remote-proxy-address local local-proxy-address;
            }
```

```
            pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
            ike-policy policy-name;
            interface-id string-value;
        }
        l2tp {
            aaa-access-profile profile-name;
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions number;
            maximum-sessions-per-tunnel number;
            multilink {
                drop-timeout milliseconds;
                fragment-threshold bytes;
            }
            override-result-code session-out-of-resource;
            ppp-authentication (chap | pap);
            ppp-profile profile-name;
            service-profile profile-name(parameter)&profile-name;
            sessions-limit-group limit-group-name;
            shared-secret shared-secret;
        }
        pap-password pap-password;
        ppp {
            cell-overhead;
            encapsulation-overhead bytes;
            framed-ip-address ip-address;
            framed-pool framed-pool;
            idle-timeout seconds;
            interface-id interface-id;
            keepalive seconds;
            primary-dns primary-dns;
            primary-wins primary-wins;
            secondary-dns secondary-dns;
            secondary-wins secondary-wins;
        }
        user-group-profile profile-name;
    }
    domain-name-server;
    domain-name-server-inet;
    domain-name-server-inet6;
    local {
        flat-file-profile profile-name;
```

```
    }
    preauthentication-order preauthentication-method;
    provisioning-order  (gx-plus | jsrc | pcrf);
    radius {
        accounting-server [ ip-address ];
        attributes {
            exclude {
                attribute-name packet-type;
                standard-attribute number {
                    packet-type [ access-request | accounting-off | accounting-on | accounting-
start | accounting-stop ];
                }
                vendor-id id-number {
                    vendor-attribute vsa-number {
                        packet-type [ access-request | accounting-off | accounting-on |
accounting-start | accounting-stop ];
                    }
                }
            }
            ignore {
                dynamic-iflset-name;
                framed-ip-netmask;
                idle-timeout;
                input-filter;
                logical-system:routing-instance;
                output-filter;
                session-timeout;
                standard-attribute number;
                vendor-id id-number {
                    vendor-attribute vsa-number;
                }
            }
        }
        authentication-server [ ip-address ];
        options {
            accounting-session-id-format (decimal | description);
            calling-station-id-delimiter delimiter-character;
            calling-station-id-format {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
                interface-text-description;
                mac-address;
```

```
            nas-identifier;
            stacked-vlan;
            vlan;
        }
        chap-challenge-in-request-authenticator;
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm  (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-channel;
            exclude-sub-interface;
        }
        juniper-access-line-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            pw-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width:
                slot-width width;
                vci-width width:
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            order {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
```

```
                    interface-text-description;
                    nas-identifier;
                    postpend-vlan-tags;
                }
                postpend-vlan-tags;
            }
            nas-port-type {
                ethernet {
                    port-type;
                }
            }
            override {
                calling-station-id remote-circuit-id;
                nas-ip-address tunnel-client-gateway-address;
                nas-port tunnel-client-nas-port;
                nas-port-type tunnel-client-nas-port-type;
            }
            remote-circuit-id-delimiter;
            remote-circuit-id-fallback {
            remote-circuit-id-format;
                agent-circuit-id;
                agent-remote-id;
            }
            revert-interval interval;
            service-activation {
                dynamic-profile (optional-at-login | required-at-login);
                extensible-service (optional-at-login | required-at-login);
            }
            vlan-nas-port-stacked-format;
        }
        preauthentication-server ip-address;
    }
    radius-server server-address {
        accounting-port port-number;
        accounting-retry number;
        accounting-timeout seconds;
        dynamic-request-port
        port port-number;
        preauthentication-port port-number;
        preauthentication-secret password;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
```

```
        max-outstanding-requests value;
         source-address source-address;
        timeout seconds;
    }
    service {
        accounting {
            statistics (time | volume-time);
            update-interval minutes;
        }
        accounting-order (activation-protocol | local | radius);
    }
    session-limit-per-username number;
    session-options {
        client-idle-timeout minutes;
        client-idle-timeout-ingress-only;
        client-session-timeoutminutes;
        pcc-context {
            input-service-filter-name filter-name;
            input-service-set-name service-set-name;
            ipv6-input-service-filter-name filter-name;
            ipv6-input-service-set-name service-set-name;
            ipv6-output-service-filter-name filter-name;
            ipv6-output-service-set-name service-set-name;
            output-service-filter-name filter-name;
            output-service-set-name service-set-name;
            profile-name pcef-profile-name;
        }
        strip-user-name {
            delimiter [ delimiter ];
            parse-direction (left-to-right | right-to-left);
        }
    }
    subscriber username {
        delegated-pool delegated-pool-name;
        framed-ip-address ipv4-address;
        framed-ipv6-pool ipv6-pool-name;
        framed-pool ipv4-pool-name;
        password password;
        target-logical-system logical-system-name <target-routing-instance (default | routing-
instance-name>;
        target-routing-instance (default | routing-instance-name);
```

```
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a subscriber access profile that includes subscriber access, L2TP, or PPP properties.

Options

profile-name—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Point-to-Point Protocol (PPP)*
- *Layer 2 Tunneling Protocol (L2TP)*
- *L2TP LNS Inline Service Interfaces*
- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol*

JSRC for Subscriber Provisioning and Accounting | 1467

Configuring Service Accounting in Local Flat Files

AAA Service Framework Overview | 2

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

process-inform

IN THIS SECTION

- Syntax | 2213
- Hierarchy Level | 2213
- Description | 2214
- Default | 2215
- Required Privilege Level | 2215
- Release Information | 2215

Syntax

```
process-inform {
    pool pool-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name overrides],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name interface interface-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
interface interface-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name interface interface-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
interface interface-name overrides],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
```

Description

Enable the processing of DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) sent from the client to request DHCP options. For DHCP local servers, the messages are also passed to the configured server list.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

Information request messages are not processed.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Enabling Processing of Client Information Requests | 777

Overriding the Default DHCP Local Server Configuration Settings | 707

protocol-master

IN THIS SECTION

- Syntax | 2215
- Hierarchy Level | 2216
- Description | 2217
- Options | 2217
- Required Privilege Level | 2217
- Release Information | 2217

Syntax

```
protocol-master (inet | inet6);
```

Hierarchy Level

[edit bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *group-name* dual-stack-group *dual-stack-group-name*],
[edit bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* forwarding-options dhcp-relay dual-stack-group],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* system services dhcp-local-server dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* system services dhcp-local-server dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],

```
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name],
[edit routing-instances name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name],
[edit vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
```

Description

Select family as protocol primary. In some customer use cases, CPE devices have reachability to multiple BNG routers for load sharing. The CPE's DHCP client broadcasts a DHCP protocol handshake to the reachable routers; more than one router may respond. In a DHCP dual-stack environment, the DHCPv4 and DHCPv6 protocol handshakes are independent of each other, meaning that each arm of the subscriber session could connect to a different router. You can avoid this situation by specify the primary protocol. For a given dual-stack subscriber, this configuration causes the rejection of any binding attempt from the secondary address family client when a binding is not currently active for the primary protocol family. If bindings are currently active for both arms when the primary protocol family binding is released or deleted, then the binding for the secondary address family is also torn down.

Options

inet INET family has protocol primary behavior.

inet6 INET6 family has protocol primary behavior.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

- Configuring Access Profile Options for Interactions with RADIUS Servers | 483
- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces | 703
- *Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

protocols (Dynamic Profiles)

IN THIS SECTION

- Syntax | 2218
- Hierarchy Level | 2220
- Description | 2220
- Default | 2221
- Required Privilege Level | 2221
- Release Information | 2221

Syntax

```
protocols {
    igmp {
        interface interface-name {
            accounting;
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave
            log-interval seconds;
            no-accounting;
            oif-map;
            passive;
```

```
            promiscuous-mode;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name
            static {
                group group {
                    source source;
                }
            }
            version version;
        }
    }
    mld {
         interfaceinterface-name {
            (accounting | no-accounting);
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave;
            log-interval seconds;
            oif-map;
            passive;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name;
            static {
                group multicast-group-address {
                    exclude;
                    group-count number;
                    group-increment increment;
                    source ip-address {
                        source-count number;
                        source-increment increment;
                    }
                }
            }
            version version;
        }
    }
    router-advertisement {
        interfaceinterface-name {
             current-hop-limitnumber;
            default-lifetime seconds;
            dns-server-address
```

```
            (managed-configuration | no-managed-configuration);
            max-advertisement-interval seconds;
            min-advertisement-interval seconds;
            (other-stateful-configuration | no-other-stateful-configuration);
            prefixprefix {
                (autonomous | no-autonomous);
                (on-link | no-on-link);
                preferred-lifetime seconds;
                valid-lifetime seconds;
            }
            reachable-time milliseconds;
            retransmit-timer milliseconds;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Enable Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) Protocol on the router and configure interface-specific values on dynamic interfaces for each protocol.

PIM and MLD manage multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use these protocols to learn which groups have members on each of their attached physical networks. Enable IGMP for the router to receive IPv4 or IPv6 multicast traffic. Enable MLD for the router to receive IPv6 multicast traffic. MLD is needed only for IPv6 networks.

You can also use this statement to enable router advertisement for IPv6 Neighbor Discovery protocol and configure interface-specific values on dynamic interfaces.

You can configure these protocols in a dynamic client profile or a dynamic service 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.

Default

IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit dynamic-profiles profile-name protocols mld]` and `[edit dynamic-profiles profile-name protocols router-advertisement]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- Configuring IGMP
- Configuring MLD

provisioning-order (Diameter Applications)

IN THIS SECTION

- Syntax | 2222
- Hierarchy Level | 2222
- Description | 2222
- Options | 2222
- Required Privilege Level | 2222
- Release Information | 2222

Syntax

```
provisioning-order (gx-plus | jsrc | pcrf);
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure AAA to use the specified application for subscriber service provisioning.

Options

gx-plus—Specify Gx-Plus as the application used to communicate with a PCRF server for subscriber service provisioning. Sets the Subscription-Id-Type Diameter AVP sub-attribute (450) to 4 (END_USER_PRIVATE) and sets the Subscription-Id-Data Diameter AVP sub-attribute (444) to `reserved`. Both of these sub-attributes are conveyed in the Diameter AVP Subscription-ID (443) by a CCR-I message.

`jsrc`—Specify JSRC as the application used to communicate with the SAE for subscriber service provisioning. JSRC is used in an SRC environment to request services from the SAE for an authenticated subscriber. JSRC attempts to activate these services. If successful, JSRC returns an ACK message. If unsuccessful, the subscriber is denied access.

pcrf—Specify Policy Control and Charging Rules Function (PCRF) as the application used to request provisioning from the PCRF server over the Gx protocol. If you change this configuration, any existing subscriber sessions are unaffected.

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 9.6.

Support for Gx-Plus introduced in Junos OS Release 11.2.

`pcrf` option added in Junos OS Release 16.2.

RELATED DOCUMENTATION

- JSRC Configuration Overview | **1476**
- Provisioning Subscribers with JSRC | **1478**
- Configuring Gx-Plus | **1404**
- Provisioning Subscribers with Gx-Plus | **1407**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

proxy-mode

IN THIS SECTION

- Syntax | **2223**
- Hierarchy Level | **2223**
- Description | **2224**
- Required Privilege Level | **2224**
- Release Information | **2224**

Syntax

```
proxy-mode;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name  overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay  overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
```

```
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay  overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name  overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay  overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.

You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- DHCP Relay Proxy Overview | **698**
- Extended DHCP Relay Agent Overview | **695**
- Enabling DHCP Relay Proxy Mode | **716**

qos-adjust

IN THIS SECTION

- For Junos OS Release 17.3 and Earlier Releases | 2225
- For Junos OS Release 17.4 and Later Releases | 2225
- Hierarchy Level | 2226
- Description | 2226
- Required Privilege Level | 2226
- Release Information | 2226

For Junos OS Release 17.3 and Earlier Releases

```
qos-adjust {
    adsl-bytes bytes;
    adsl2-bytes bytes;
    adsl2-plus-bytes bytes;
    other-bytes;
    other-overhead-adjust;
    sdsl-bytes bytes;
    sdsl-overhead-adjust percentage;
    vdsl-bytes bytes;
    vdsl-overhead-adjust percentage;
    vdsl2-bytes bytes;
    vdsl2-overhead-adjust percentage;
}
```

For Junos OS Release 17.4 and Later Releases

```
qos-adjust;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Specify that the ANCP agent reports data rates for downstream traffic to CoS. When this statement is not configured, the ANCP agent does not report traffic rates to CoS.

For Junos OS Release 17.3 and earlier releases, configure the values by which the actual downstream data rates are adjusted. The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Starting in Junos OS Release 17.4, all of the previously supported subordinate statements are deprecated. The ANCP agent ignores them if they are present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustments with the `access-line` statement at the `[edit system]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Subordinate statements deprecated in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**

Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**

Configuring the ANCP Agent | **1253**

qos-adjust-adsl

IN THIS SECTION

- Syntax | 2227
- Hierarchy Level | 2227
- Description | 2227
- Options | 2228
- Required Privilege Level | 2228
- Release Information | 2228

Syntax

```
qos-adjust-adsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

qos-adjust-adsl2

IN THIS SECTION

Syntax | 2229

- Hierarchy Level | 2229
- Description | 2229
- Options | 2229
- Required Privilege Level | 2230
- Release Information | 2230

Syntax

```
qos-adjust-adsl2 adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL2 line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl2` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

qos-adjust-adsl2-plus

IN THIS SECTION

- Syntax | 2231
- Hierarchy Level | 2231
- Description | 2231
- Options | 2231
- Required Privilege Level | 2231
- Release Information | 2231

Syntax

```
qos-adjust-adsl2-plus adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL2+ line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl2-plus` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306
- Configuring the ANCP Agent | 1253

qos-adjust-other

IN THIS SECTION

- Syntax | 2232
- Hierarchy Level | 2233
- Description | 2233
- Options | 2233
- Required Privilege Level | 2233
- Release Information | 2233

Syntax

```
qos-adjust-other adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an access line of DSL type OTHER. The ANCP agent reports the adjusted rate only to AAA.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-other` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

qos-adjust-sdsl

IN THIS SECTION

- Syntax | 2234
- Hierarchy Level | 2234
- Description | 2235
- Options | 2235
- Required Privilege Level | 2235
- Release Information | 2235

Syntax

```
qos-adjust-sdsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an SDSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-sdsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

qos-adjust-vdsl

IN THIS SECTION

- Syntax | 2236
- Hierarchy Level | 2236
- Description | 2236
- Options | 2237
- Required Privilege Level | 2237
- Release Information | 2237

Syntax

```
qos-adjust-vdsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on a VDSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-vdsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

qos-adjust-vdsl2

IN THIS SECTION

- Syntax | 2238
- Hierarchy Level | 2238
- Description | 2238
- Options | 2239
- Required Privilege Level | 2239
- Release Information | 2239

Syntax

```
qos-adjust-vdsl2 adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on a VDSL2 line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-vdsl2` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates | 1306

Configuring the ANCP Agent | 1253

radius (Access Profile)

IN THIS SECTION

Syntax | 2240

- Hierarchy Level | 2243
- Description | 2243
- Options | 2243
- Required Privilege Level | 2243
- Release Information | 2244

Syntax

```
radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
            attribute-name packet-type;
            standard-attribute number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
| accounting-stop ];
            }
            vendor-id id-number {
                vendor-attribute vsa-number {
                    packet-type [ access-request | accounting-off | accounting-on | accounting-
start | accounting-stop ];
                }
            }
        }
        ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            idle-timeout;
            input-filter;
            logical-system-routing-instance;
            output-filter;
            session-timeout;
            standard-attribute number;
            vendor-id id-number {
                vendor-attribute vsa-number;
            }
        }
    }
```

```
    authentication-server [ ip-address ];
    options {
        accounting-session-id-format (decimal | description);
        calling-station-id-delimiter delimiter-character;
        calling-station-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        chap-challenge-in-request-authenticator;
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm  (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-channel;
            exclude-sub-interface;
        }
        ip-address-change-notify message;
        juniper-access-line-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width:
                slot-width width;
                vci-width width:
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
```

```
            interface-text-description;
            nas-identifier;
            order {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
                interface-text-description;
                nas-identifier;
                postpend-vlan-tags;
            }
            postpend-vlan-tags;
        }
        nas-port-type {
            ethernet {
                port-type;
            }
        }
        override {
            calling-station-id remote-circuit-id;
            nas-ip-address tunnel-client-gateway-address;
            nas-port tunnel-client-nas-port;
            nas-port-type tunnel-client-nas-port-type;
        }
        remote-circuit-id-delimiter;
        remote-circuit-id-fallback;
        remote-circuit-id-format {
            agent-circuit-id;
            agent-remote-id;
        }
        revert-interval interval;
        service-activation {
            dynamic-profile (optional-at-login | required-at-login);
            extensible-service (optional-at-login | required-at-login);
        }
        vlan-nas-port-stacked-format;
    }
    preauthentication-server ip-address;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

Options

accounting-server (MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

- **Values:** *ip-address*—IP version 4 (IPv4) address.

authentication-server (SRX Series only) Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

- **Values:** *ip-address*—IPv4 address.

preauthentication-server (MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

> **NOTE**: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the `exclude` statement.

- **Values:** *ip-address*—IPv4 address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration | 550

RADIUS Logical Line Identification | 543

radius-disconnect (DHCP Local Server)

IN THIS SECTION

- Syntax | 2244
- Hierarchy Level | 2244
- Description | 2245
- Default | 2245
- Required Privilege Level | 2245
- Release Information | 2246

Syntax

```
radius-disconnect;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure trigger],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure trigger],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure trigger],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure
trigger],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure
trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure trigger],
[edit system services dhcp-local-server reconfigure trigger],
[edit system services dhcp-local-server dhcpv6 reconfigure trigger],
[edit system services dhcp-local-server group group-name reconfigure trigger],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure trigger]
```

Description

Configure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces to be reconfigured when a RADIUS-initiated disconnect is received by the DHCP client or group of clients. A group configuration takes precedence over a DHCP local server configuration.

Default

The client is deleted when a RADIUS-initiated disconnect is received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | 876

radius-flow-tap

IN THIS SECTION

- Syntax | 2246
- Hierarchy Level | 2247
- Description | 2248
- Options | 2248
- Required Privilege Level | 2249
- Release Information | 2249

Syntax

```
radius-flow-tap {
    forwarding-class class-name;
    interfaces interface-name;
    logical-system logical-system-name name <routing-instance routing-instance>;
    multicast-interception;
    policy policy-name {
        inet
{
            drop-policyrule-name {
                from
```

```
{
                    apply-groups group-name;
                    apply-groups-except group-name;
                    destination-address address;
                    destination-port port-number;
                    dscp dscp-value;
                    protocol protocol;
                    source-address address;
                    source-port port-number;
                }
            }
        }
        inet6 {
            drop-policy rule-name {
                from {
                    apply-groups group-name;
                    apply-groups-except group-name;
                    destination-address address;
                    destination-port port-number;
                    dscp dscp-value;
                    protocol protocol;
                    source-address address;
                    source-port port-number;
                }
            }
        }
    }
    snmp (
        notify-targets ip-address;
    }
    routing-instance routing-instance-name;
    source-ipv4-address ipv4-address;
)
```

Hierarchy Level

```
[edit services]
```

Description

Configure the radius-flow-tap service for subscriber secure policy mirroring. Both RADIUS-initiated and Dynamic Tasking Control Protocol (DTCP)-initiated mirroring are supported.

Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service. The FlowTapLite service is a version of the flow-tap service (`[edit services flow-tap]`) that is configured only on tunnel interfaces on MX Series routers.

In earlier releases, the radius-flow-tap and FlowTapLite services cannot run concurrently on an MX Series router, preventing you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

forwarding-class *class-name*
Specify the forwarding class that is applied to mirrored packets sent to a mediation device.

interfaces *interface-name*
Assign virtual tunnel interfaces to mirror the interfaces created by extensible subscriber services manager (ESSM)

If a currently used tunnel interface is deleted from the pool of interfaces, the active mirroring sessions are redistributed from the deleted interface to other tunnel interfaces in the pool. Also, when a new tunnel interface is added into the pool, the service adds the new interface to the list of interfaces available for new mirroring sessions or for existing sessions transferred from a failed interface.

logical-system *logical-system-name*
Specify the logical system that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring. When you specify a logical system, you must also specify a routing instance.

- **Default:** Logical system `default`

multicast-interception
Enables subscriber secure policy to mirror IPv4 multicast traffic sent to subscribers. It enables the mirroring of multicast traffic for all subscribers on the chassis.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

routing-instance *routing-instance-name*
Specify the routing instance that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring.

- **Default:** Routing instance `default`

snmp notify-targets *ip-address* Specify the IP address for a target mediation device (trap target) to receive SNMPv3 encrypted trap notifications subscriber secure policy mirroring trap. Only these configured targets can receive the notifications. This is required for secure SNMPv3 notifications for subscriber secure policy mirroring. If you configure multiple targets, you must configure them one at a time.

source-ipv4-address *ipv4-address* Specify the source IPv4 address used in the IP header that is prepended to mirrored packets sent to a mediation device.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`logical-system` option added in Junos OS Release 15.1R3 for enhanced subscriber management.

`multicast-interception` option added in Junos OS Release 11.4.

`snmp notify-targets` option added in Junos OS Release 16.1R1.

`routing-instance` option added in Junos OS Release 15.1R3 for enhanced subscriber management.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service.

RELATED DOCUMENTATION

Configuring Support for Subscriber Secure Policy Mirroring

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring

Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview

Subscriber Secure Policy Overview

radius-options (Access)

IN THIS SECTION

- Syntax | 2250
- Hierarchy Level | 2250
- Description | 2251
- Required Privilege Level | 2251
- Release Information | 2251

Syntax

```
radius-options {
    interim-rate interim-rate;
    request-rate rate;
    revert-interval seconds;
    timeout-grace seconds;
    unique-nas-port {
        chassis-id chassis-id;
        chassis-id-width chassis-id-width;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure RADIUS options that apply to all RADIUS servers globally.

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 8.3.

RELATED DOCUMENTATION

- RADIUS Servers and Parameters for Subscriber Access | 476
- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | 482
- Configuring Authentication and Accounting Parameters for Subscriber Access | 550

radius-options (Interfaces)

IN THIS SECTION

- Syntax | 2252
- Hierarchy Level | 2252
- Description | 2252
- Required Privilege Level | 2252
- Release Information | 2253

Syntax

```
radius-options {
    nas-port-options nas-port-options-name {
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked;
            stacked-vlan-width width;
            vci-width width:
            vlan-width width;
            vpi-width width;
        }
        nas-port-type port-type;
        stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
        vlan-ranges (any | low-tag-high-tag);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure RADIUS options to set 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 remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 527

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | 526

radius-server

IN THIS SECTION

- Syntax | 2253
- Hierarchy Level | 2254
- Description | 2254
- Options | 2254
- Required Privilege Level | 2258
- Release Information | 2258

Syntax

```
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    max-outstanding-requests value;
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
```

```
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name]
```

Description

Configure RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple `radius-server` statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address IPv4 or IPv6 address of the RADIUS server.

accounting-port Configure the port number on which to contact the RADIUS accounting server.

> **NOTE**: Specifying the accounting port is optional, and port 1813 is the default. However, we recommend that you configure it in order to avoid confusion, as some RADIUS servers might refer to an older default.

- **Values:** *port-number*—Port number on which to contact the RADIUS accounting server. Most RADIUS servers use port 1813, as specified in RFC 2866.
- **Default:** 1813

accounting-retry Configure the number of times the device retransmits RADIUS accounting messages when no response is received from the server. When you do not configure this statement, the number of retry attempts is determined by the `retry` statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

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.

- **Values:** *number*—Number of retry attempts.
- **Range:** 0 through 100
- **Default:** 0 (disabled)

accounting-timeout Configure how long the local device waits to receive a response from a RADIUS accounting server before retransmitting the message. When you do not configure this statement, the length of the timeout is determined by the `timeout` statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

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.

- **Values:** *seconds*—Duration of timeout period.
- **Range:** 0 through 1000 seconds
- **Default:** 0 (disabled)

dynamic-request-port

Specify the port that the router monitors for dynamic (CoA) requests from the specified RADIUS servers. You can configure a port globally or for a specific access profile.

You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

> **NOTE**: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

- **Values:** *port-number*—Number of the monitored port.
- **Default:** 3799 (as specified in RFC 5176)

max-outstanding-requests

Configure the maximum number of outstanding requests for this RADIUS server. An increase in this value is immediate while a decrease is more gradual if the current number of outstanding requests exceeds the new value.

- **Values:** *requests*—Maximum number of outstanding requests for this RADIUS server.
- **Range:** 0 through 2000 outstanding requests per server
- **Default:** 1000 outstanding requests per server

port

Configure the port number on which to contact the RADIUS server.

- **Values:** *port-number*—Port number on which to contact the RADIUS server.
- **Default:** 1812 (as specified in RFC 2865)

preauthentication-port

Configure the port number on which to contact the RADIUS server for logical line identification (LLID) preauthentication requests. If you do not configure a separate UDP port for preauthentication purposes, the same UDP port that you configure for authentication messages by including the `port` *`port-number`* statement is used.

- **Values:** *port-number*—Port number used for preauthentication requests to contact the RADIUS server.

preauthentication-secret

Configure the password to use with the RADIUS server for LLID preauthentication requests. If you do not configure a separate UDP password for

preauthentication purposes, the same password that you configure for authentication messages by including the `secret` *`password`* statement is used. The secret password used by the local router must match that used by the server.

- **Values:** *password*—Password to use. To include spaces enclose the character string in quotation marks.

retry

Specify the number of times that the device is allowed to attempt to contact a RADIUS authentication or accounting server. You can override the retry limit for accounting servers with the `accounting-retry` statement.

> **NOTE**: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **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.

- **Values:** *attempts*—Number of times that the router is allowed to attempt to contact a RADIUS server.
- **Range:** 1 through 100
- **Default:** 3

routing-instance

Configure the routing instance used to send RADIUS packets to the RADIUS server.

- **Values:** *routing-instance-name*—Routing instance name.

source-address

Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. Support for IPv6 *`source-address`* was introduced in Junos OS Release 16.1.

- **Values:** *source-address*—Valid IPv4 or IPv6 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

timeout

Configure the amount of time that the local device waits to receive a response from RADIUS authentication and accounting servers. You can override the timeout value for accounting servers with the `accounting-timeout` statement.

> **NOTE**: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **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.

- **Values:** *seconds*—Amount of time to wait.
- **Range:** 1 through 1000 seconds
- **Default:** 3 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`max-outstanding-requests` introduced in Junos OS Release 11.4.

`accounting-retry` and `accounting-timeout` introduced in Junos OS Release 14.1.

`dynamic-request-port` option added in Junos OS Release 14.2R1 for MX Series routers.

`preauthentication-port` and `preauthentication-secret` options added in Junos OS Release 15.1 for MX Series routers.

`accounting-port` introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

Support for IPv6 *server-address* introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- RADIUS Authentication and Accounting Basic Configuration | **550**
- *Configuring the PPP Password Authentication Protocol*
- *Configuring RADIUS Authentication for L2TP*
- *RADIUS Authentication*
- Configuring RADIUS-Initiated Dynamic Request Support
- RADIUS Logical Line Identification | **543**
- show network-access aaa statistics | **2958**
- clear network-access aaa statistics | **2602**

range (Address-Assignment Pools)

IN THIS SECTION

- Syntax | **2259**
- Hierarchy Level | **2260**
- Description | **2260**
- Options | **2260**
- Required Privilege Level | **2260**
- Release Information | **2260**

Syntax

```
range range-name {
    high upper-limit;
```

```
    low lower-limit;
    prefix-length prefix-length;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Description

Configure a named range of IPv4 addresses or IPv6 prefixes, used within an address-assignment pool.

Options

high *upper-limit*—Upper limit of an address range or IPv6 prefix range.

low *lower-limit*—Lower limit of an address range or IPv6 prefix range.

prefix-length *prefix-length*—Assigned length of the IPv6 prefix.

range-name—Name assigned to the range of IPv4 addresses or IPv6 prefixes.

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 9.0.

IPv6 support introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management | 1138

rapid-commit (DHCPv6 Local Server)

IN THIS SECTION

- Syntax | 2261
- Hierarchy Level | 2261
- Description | 2261
- Default | 2262
- Required Privilege Level | 2262
- Release Information | 2262

Syntax

```
rapid-commit;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 ...],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 ...]
```

Description

Configure DHCPv6 local server to recognize the Rapid Commit option (DHCPv6 option 14) in DHCPv6 solicit messages sent from the DHCPv6 client. When rapid commit is enabled for both DHCPv6 local server and the DHCPv6 client, a two-message handshake is used instead of the standard four-message handshake. You can enable rapid commit support on DHCPv6 local server globally, for a named group, or for a specific interface.

Default

Rapid commit support is not enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring DHCPv6 Rapid Commit (MX Series, EX Series) | 912

Overriding the Default DHCP Local Server Configuration Settings | 707

reachable-time (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2262
- Hierarchy Level | 2263
- Description | 2263
- Options | 2263
- Required Privilege Level | 2263
- Release Information | 2263

Syntax

```
reachable-time milliseconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Set the length of time that a node considers a neighbor reachable until another reachability confirmation is received from that neighbor.

Options

milliseconds—Reachability time limit.

- **Range:** 0 through 3,600,000 milliseconds
- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

realm-delimiter (Domain Map)

IN THIS SECTION

- Syntax | 2264
- Hierarchy Level | 2264
- Description | 2264
- Default | 2264
- Options | 2264
- Required Privilege Level | 2265
- Release Information | 2265

Syntax

```
realm-delimiter [delimiter-character];
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the characters that the router uses to separate usernames from realm names.

Default

none

Options

delimiter-character—One or more characters used as delimiters. You can specify a maximum of eight delimiters. You cannot use the semicolon (;) as a delimiter. Do not include spaces between characters.

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 14.1.

RELATED DOCUMENTATION

Specifying Domain and Realm Name Delimiters | 668

Configuring Domain and Realm Name Usage for Domain Maps | 667

realm-parse-direction (Domain Map)

IN THIS SECTION

- Syntax | 2265
- Hierarchy Level | 2266
- Description | 2266
- Default | 2266
- Options | 2266
- Required Privilege Level | 2266
- Release Information | 2266

Syntax

```
realm-parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the direction in which the router searches for the realm name in a username.

Default

`left-to-right`

Options

`left-to-right`—The router searches starting at the left-most character. When the router reaches a realm delimiter, it uses anything to the left of the delimiter as the realm name.

`right-to-left`—The router searches starting at the right-most character. When the router reaches a realm delimiter, it uses anything to the left of the realm as the domain name.

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 14.1.

RELATED DOCUMENTATION

Specifying the Parsing Direction for Domain and Realm Names | 670

Configuring Domain and Realm Name Usage for Domain Maps | 667

reauthenticate (DHCP Local Server)

IN THIS SECTION

- Syntax | 2267
- Hierarchy Level | 2267
- Description | 2268
- Options | 2269
- Required Privilege Level | 2269
- Release Information | 2269

Syntax

```
reauthenticate (<lease-renewal> <remote-id-mismatch > <actual-data-rate-change>);
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server],
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-
group name],
[edit logical-systems name system services dhcp-local-server],
[edit logical-systems name system services dhcp-local-server dhcpv6],
[edit logical-systems name system services dhcp-local-server dual-stack-group name],
[edit routing-instances name system services dhcp-local-server ],
[edit routing-instances name system services dhcp-local-server dhcpv6],
[edit routing-instances name system services dhcp-local-server
dual-stack-group name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dual-stack-group name]
```

Description

Enable DHCP client reauthentication, that is, trigger jdhcpd to request reauthentication from authd, which in turn reissues the RADIUS Access-Request for subscriber authentication. The purpose of the reauthentication is to change characteristics of the subscriber session, such as activating subscriber services or changing attributes. You can use reauthentication as an alternative to a RADIUS CoA request.

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 specify that reauthentication occurs in response to all DHCP renew, rebind, discover, or solicit messages or only in response to discover and solicit messages that include a new (different) Agent Remote ID for the DHCP client.

You can use the Juniper Networks VSA, Reauthentication-On-Renew (26-206) as an alternative to the CLI configuration to enable reauthentication. The `reauthenticate` statement overrides the VSA when the VSA is present with a value of disable.

NOTE: Reauthentication for dual-stack, single-session subscribers requires that the `on-demand-address-allocation` statement is configured for the dual-stack group. This is true whether you enable reauthentication with the `reauthenticate` statement or the Reauthenticate-On-Renew VSA (26-206).

NOTE: You cannot configure both the `reauthenticate` statement and the remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) 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 (DHCP Local Server and DHCP Relay Agent) for DHCPv6 at the `[edit system services dhcp-local-server dhcpv6]` or for a specific group at the `[edit system services dhcp-local-server group` *`name`*`]` hierarchy level, and so on.

NOTE: Reauthentication does not support Extensible Services Subscriber Management (essmd) services. Activation or deactivation of any such service causes the request to fail.

Options

lease-renewal Reauthenticate when a renew, rebind, discover, or solicit message is received from the DHCP client. This re-authentication is an alternative to RADIUS Change of Authorization (CoA) to change subscriber session characteristics based on the actual data rate change without interrupting the subscriber service.

remote-id-mismatch Reauthenticate when a discover or solicit message is received from the DHCP client with a new value for the DHCP client's Agent Remote ID. The change in value corresponds to a change in subscriber service plan. The Agent Remote ID is conveyed in option 82, suboption 2 for DHCPv4 clients and in option 37 for DHCPv6 clients.

actual-data-rate-change Optical line terminal (OLT) adds `option 82` with `sub-option 9` with Broadband Forum (Vendor ID 3561). It contains the sub-attributes `Actual-Data-Rate-Upstream` and `Actual-Data-Down-Stream` encoded. The decoded upstream and downstream data send to RADIUS server as a part of authentication request. RADIUS server analyzes the data received. Based on the `Actual-Data-Rate-Upstream` and `Actual-Data-Down-Stream` values, the RADIUS server selects the service profile for the subscriber interface.

When the actual data rate changes, the DHCP server re-authenticates the subscriber service. This re-authentication is an alternative to RADIUS Change of Authorization (CoA) to change subscriber session characteristics based on the actual data rate change without interrupting the subscriber service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

Support at the `[edit ... system services dhcp-local-server dual-stack-group]` hierarchy level introduced in Junos OS Release 18.1R1.

`actual-data-rate-change` option introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Configuring RADIUS Reauthentication for DHCP Subscribers | 567

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers | 556

reconfigure (DHCP Local Server)

IN THIS SECTION

- Syntax | 2270
- Hierarchy Level | 2270
- Description | 2271
- Options | 2271
- Required Privilege Level | 2271
- Release Information | 2272

Syntax

```
reconfigure {
    attempts attempt-count;
    clear-on-terminate;
    strict;
    support-option-pd-exclude;
    timeout timeout-value;
    token token-value;
     trigger {
        radius-disconnect;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-
name],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Enable dynamic reconfiguration triggered by the DHCP local server of all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. A group configuration takes precedence over a DHCP local server configuration. The `strict` statement is available only for DHCPv6.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

support-option-pd-exclude Request to exclude prefix option in the reconfigure message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

`support-option-pd-exclude` statement introduced in Junos OS Release 17.3 for the MX Series.

RELATED DOCUMENTATION

- Dynamic Reconfiguration of DHCP Servers and Clients

redundancy (M:N Subscriber Redundancy)

IN THIS SECTION

- Syntax | 2272
- Hierarchy Level | 2273
- Description | 2273
- Options | 2273
- Required Privilege Level | 2276
- Release Information | 2276

Syntax

```
redundancy {
    interface name {
        local-inet-address v4-address;
        local-inet6-address v6-address;
        shared-key string;
        virtual-inet-address virtual-inet-address;
        virtual-inet6-address virtual-inet6-address;
    }
    group group-name {
            interface interface-name {
```

```
                standby-mode [hot-standby | service-activation-on-failover];
            }
        }
    }
    no-advertise-routes-on-backup;
    protocol {
        pseudo-wire;
        vrrp;
    }
}
```

Hierarchy Level

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

Description

Enable M:N subscriber group redundancy. This global configuration applies to all subscribers in redundancy groups across the chassis.

For dual-stack subscribers, you must configure both virtual IPv4 and virtual IPv6 addresses on the interface.

Options

group *group-name* Name of the redundancy group.

hot-standby Specifies the backup BNG, which is ready to switchover the traffic from the primary BNG immediately without any disruption during primary BNG failure.

interface *interface-name* Identify the interface name for the logical access interface for the subscriber redundancy group covered by M:N redundancy.

Configuring the interfaces inside the group is for logical grouping only. You can group all the interfaces to one interface-group or you can have single interface-group for each redundancy interfaces.

NOTE: You must configure the names for all such interfaces on the chassis for all redundancy groups.

For VRRP redundancy, only ge and xe interfaces are supported for the underlying physical interface. Ethernet connections can be untagged, single-tagged, or double-tagged.

For pseudowire redundancy, only pseudowire interfaces are supported for the underlying physical interface.

NOTE: You cannot configure an interface for redundancy if it already has active non-DHCP subscribers. If the interface has existing DHCP subscribers, then redundancy is enabled for those subscribers.

local-inet-address *v4-address*

(Pseudowires only) Local IP address for the associated pseudowire interface. This address must match one of the access-facing GE interface addresses. It is unique per redundancy group identified by the ps*x*.0 interface. The DHCP relay agent uses this address as the giaddr for DHCP messages sent to a DHCP server. For dynamic subscriber Interfaces, this address is the same as the preferred-source address (the unnumbered address derived from the access-facing GE interface).

local-inet6-address *v6-address*

(Pseudowires only) Local IPv6 address for the associated pseudowire interface. This address must match one of the access-facing GE interface addresses. It is unique per redundancy group identified by the ps*x*.0 interface. The DHCP relay agent uses this address as the linkaddr for DHCP messages sent to a DHCP server. For dynamic subscriber Interfaces, this address is the same as the preferred-source address (the unnumbered address derived from the access-facing GE interface).

no-advertise-routes-on-backup

(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 BNG 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. In this mode, the DHCP server uses a common pool for all redundant subscribers and possibly all non-redundant subscribers. This mode increases the routing services overhead due to the number of routes being advertised for all subscribers. It increases the size of the routing table on the core side and can increase core-side convergence if a DHCP relay agent fails.

It also increases downstream traffic convergence. The `no-advertise-routes-on-backup` option reduces the number of routes advertised and the associated potential issues.

- **Default:** Advertise the access and framed routes to the backup BNG during subscriber login.

protocol

Specify a method for BNG M:N redundancy. You can configure both VRRP and pseudowires when the underlying access interfaces do not overlap between these methods.

- `pseudo-wire`—Use pseudowires for redundancy on pseudowire interfaces for L2VPN and Layer 2 circuit-based networks over IP/MPLS.
- `vrrp`—Use VRRP for redundancy on GE and XE interfaces.
- **Default:** `vrrp`

shared-key *string*

(Pseudowires only) Shared key that is used by the topology discovery process to identify a backup pseudowire interface on a peer BNG for a subscriber redundancy group. You configure the same shared-key on the local access interfaces on the primary and backup BNGs for a particular redundancy group. A topology discovery match is achieved when both the query and response messages include the pseudowire local access interface name, the same shared key, and the same transaction ID. You define the shared key as a string of up to 64 characters that you choose. You must not use the shared key for any other peer BNGs.

NOTE: If the interface already has non-DHCP subscribers, the router rejects the configuration. The configuration is accepted for existing DHCP relay subscribers.

NOTE: You cannot dynamically change a shared key after you have configured a key for a subscriber redundancy group (represented by the pseudowire interface). This means that if you change the value after committing it, the change has no effect. In order to have the change take effect, you must delete the entire redundancy group configuration and then add it back with the new key.

service-activation-on-failover

Specifies the interface mode on backup BNG, which uses less resources to forward traffic instantly on traffic switchover with basic statistics.

standby-mode	Specifies how to program the Packet Forwarding Engine with subscriber state when the interface is inactive.
virtual-inet-address *virtual-inet-address*	(VRRP) IPv4 address for the virtual router. This address is used as the gateway address for all redundancy peers associated with a particular subscriber redundancy group. You must specify the same virtual IP address that you also configure on VRRP so that it can create the virtual router for redundancy.
virtual-inet6-address *virtual-inet6-address*	(VRRP) IPv6 address for the virtual router. This address is used as the gateway address for all redundancy peers associated with a particular subscriber redundancy group. You must specify the same virtual IP address that you also configure on VRRP so that it can create the virtual router for redundancy.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

`local-inet-address`, `local-inet6-address`, `pseudo-wire`, and `shared-key` options added in Junos OS Release 20.1R1.

`service-activation-on-failover` introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization | 1205

relay-agent-interface-id (DHCP Local Server)

IN THIS SECTION

Syntax | 2277

- Hierarchy Level | 2277
- Description | 2277
- Required Privilege Level | 2278
- Release Information | 2278

Syntax

```
relay-agent-interface-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Creating Unique Usernames for DHCP Clients | 834

relay-agent-interface-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 2278
- Hierarchy Level | 2279
- Description | 2279
- Required Privilege Level | 2279
- Release Information | 2279

Syntax

```
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
```

```
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dual-stack-group dual-
stack-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dual-stack-group dual-stack-group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dual-stack-group ],
```

Description

Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 1752

Extended DHCP Relay Agent Overview | 695

DHCPv6 Relay Agent Overview | 916

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | 919

relay-agent-interface-id (DHCPv6 Relay Agent Username)

IN THIS SECTION

- Syntax | 2280
- Hierarchy Level | 2280
- Description | 2281
- Required Privilege Level | 2281
- Release Information | 2281

Syntax

```
relay-agent-interface-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview | 916

Creating Unique Usernames for DHCP Clients | 834

relay-agent-remote-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 2282
- Hierarchy Level | 2282
- Description | 2282
- Required Privilege Level | 2282
- Release Information | 2283

Syntax

```
relay-agent-remote-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

For MX Series routers only, `enterprise-id` and `remote-id` options introduced in Junos OS Release 12.3R3.

For MX Series routers only, the `enterprise-id` and `remote-id` options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

relay-agent-remote-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 2283
- Hierarchy Level | 2284
- Description | 2284
- Required Privilege Level | 2284
- Release Information | 2284

Syntax

```
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ... ],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ... ],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ... ]
```

Description

Specify that the DHCPv6 relay agent include Relay Agent Remote-ID (option 37) in DHCPv6 packets destined for a DHCPv6 server. Optionally specify that the option includes a prefix, the interface textual description, or both. In dual-stack environments, you can also specify that the DHCPv6 relay agent use the DHCPv4 option 82 value for DHCPv6 option 37.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Release History Table

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `relay-agent-remote-id` CLI statement in a stateless DHCPv6 relay configuration. You can configure stateless DHCPv6 relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay dhcpv6`] hierarchy level.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

DHCPv6 Relay Agent Overview | 916
DHCPv6 Relay Agent Options | 917
Configuring DHCPv6 Relay Agent Options | 917

relay-agent-remote-id (DHCPv6 Relay Agent Username)

IN THIS SECTION

- Syntax | 2285
- Hierarchy Level | 2285
- Description | 2286
- Required Privilege Level | 2286
- Release Information | 2286

Syntax

```
relay-agent-remote-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

For MX Series routers only, `enterprise-id` and `remote-id` options introduced in Junos OS Release 12.3R3.

For MX Series routers only, the `enterprise-id` and `remote-id` options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview | 916

Creating Unique Usernames for DHCP Clients | 834

relay-agent-subscriber-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 2287
- Hierarchy Level | 2287
- Description | 2288
- Required Privilege Level | 2288
- Release Information | 2288

Syntax

```
relay-agent-subscriber-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6  authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Subscriber-ID option (option 38) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

relay-agent-subscriber-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 2288
- Hierarchy Level | 2289
- Description | 2289
- Required Privilege Level | 2289
- Release Information | 2289

Syntax

```
relay-agent-subscriber-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Subscriber-ID option (option 38) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview | 916

Creating Unique Usernames for DHCP Clients | 834

relay-option (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2290
- Hierarchy Level | 2291
- Description | 2291
- Required Privilege Level | 2291
- Release Information | 2291

Syntax

```
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name 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 ...]
```

Description

Configure the extended DHCP relay agent selective processing that is based on DHCP options in DHCP client packets and specify the action to perform on client traffic. You can configure support globally or for a named group of interfaces, and for either DHCP or DHCPv6 relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

relay-option-vendor-specific (dhcpv6)

IN THIS SECTION

- Syntax | 2292
- Hierarchy Level | 2292
- Description | 2292
- Required Privilege Level | 2293
- Release Information | 2293

Syntax

```
relay-option-vendor-specific {
    host-name;
    location;
}
```

Hierarchy Level

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

Description

Add vendor-specific information to the vendor-specific option (17) field of DHCPv6 control messages on server-facing interfaces. The vendor-specific information can be a hostname, a location (such as a unique connection identifier), or both. The hostname can be a string of characters such as **Juniper-AB-1**. The location should be specified as interface, VLAN ID, and if applicable, stacked VLAN ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243. The enterprise ID is 2636. The hostname is option-data 1, and the location is option-data 2. The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query a single entity to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

relay-option-82

IN THIS SECTION

- Syntax | 2293
- Syntax (QFX Series) | 2294
- Hierarchy Level | 2295
- Description | 2295
- Required Privilege Level | 2296
- Release Information | 2296

Syntax

```
relay-option-82 {
    circuit-id {
        include-irb-and-l2;
        keep-incoming-circuit-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
```

```
        use-vlan-id;
    }
    remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    server-id-override;
    vendor-specific{
        host-name;
        location;
    }
}
```

Syntax (QFX Series)

```
relay-option-82 {
    circuit-id {
        include-irb-and-l2;
        keep-incoming-circuit-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    server-id-override;
    link-selection;
    vendor-specific{
        host-name;
        location;
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

To enable insertion of option 82 information in DHCP packets, you must specify at least one of the `circuit-id` or `remote-id` statements.

You can use the `relay-option-82` statement and its subordinate statements at the `[edit forwarding-options dhcp-relay]` hierarchy level to control insertion of option 82 information globally, or at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level to control insertion of option 82 information for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the `delete relay-option-82` statement.

Starting in Junos OS Release 21.2R1, on QFX Series devices, we've introduced `link-selection` statement at the `edit forwarding-options dhcp-relay relay-option-82` hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`link-selection` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

dhcp-relay | 1752

relay-server-group (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 2296
- Hierarchy Level | 2297
- Description | 2297
- Options | 2297
- Required Privilege Level | 2297
- Release Information | 2297

Syntax

```
relay-server-group relay-server-group;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with),
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with),
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with),
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with),
[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 ...]
```

Description

Relay DHCP client packets to the specified group of DHCP servers when you use the DHCP relay selective processing feature. You can configure the relay operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.

Options

relay-server-group	Name of DHCP server group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

relay-source

IN THIS SECTION

- Syntax | 2298
- Hierarchy Level | 2298
- Description | 2299
- Options | 2299
- Required Privilege Level | 2299
- Release Information | 2299

Syntax

```
relay-source interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
interface interface-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name interface interface-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Configure DHCP relay server to use the configured loopback address as the source address in both the IP header and in DHCP messages relayed to the server.

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 BNG firewall. In that case, DHCP unicast packets do not pass through the firewall and are discarded. This configuration statement places the loopback address in IP headers and DHCP messages, which enables DHCP unicast packets to pass through the firewall to the DHCP server.

Options

interface-name	Specify the loopback interface. The *interface-name* must be `lo0`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall | **717**

remote-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2300
- Hierarchy Level | 2300
- Description | 2301
- Required Privilege Level | 2302
- Release Information | 2302

Syntax

```
remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82]
```

Description

Specify the Agent Remote ID suboption (suboption 2) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix, textual description, or VLAN tag.

NOTE: For Layer 3 interfaces, when you configure relay-option-82 only, the Agent Remote ID is the default If no VLAN tags are configured, then the default is the logical interface device (IFL) name. For integrated routing and bridging (IRB) interfaces, the default is the Layer 2 IFL name and bridge domain name.

The interface to bridge domain relationship may be implicit (the interface is mapped to the bridge domain by the system based on 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.

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs), stacked VLANs (S-VLANs), or bridge domains is as follows:

```
(fe | ge)-fpc/pic/port.subunit
```

NOTE: For remote systems, the *subunit* is required and is used to differentiate an interface.

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:

```
(fe | ge)-fpc/pic/port:vlan-id
```

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:

```
(fe | ge)-fpc/pic/port:svlan-id--vlan-id
```

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
```

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Release History Table

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `remote-id` CLI statement in a stateless DHCP relay configuration. You can configure stateless DHCP relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay`] hierarchy level.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | 751

Using DHCP Relay Agent Option 82 Information | 751

remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2303
- Hierarchy Level | 2303
- Description | 2304
- Required Privilege Level | 2305
- Release Information | 2305

Syntax

```
remote-id-mismatch disconnect;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
```

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Configure the DHCP local server or DHCP relay agent to detect a mismatch in the Agent Remote ID value to trigger a new connection request. Information about a subscriber's service plan is encoded in the Agent Remote ID, which is conveyed in option 82, suboption 2, for DHCPv4 clients and in option 37 for DHCPv6 clients. When a subscriber session is activated, the Agent Remote ID value for the authorized service plan is stored in the session database. When you configure `remote-id-mismatch`, the DHCP local server and relay agent inspect incoming renew and rebind messages and compare the Agent Remote ID in the message against the initial value that DHCP stored in the database. When DHCP local server discovers a mismatch between the stored value and the value in the message, DHCP local server sends a NAK to the client and tears down the client binding. If the client is a DHCPv6 client, because DHCPv6 does not support an explicit NAK message, the local server sends a reply packet with lifetime set to 0 to signify a logical NAK.

When DHCP relay agent discovers the mismatch, it sends a NAK or logical NAK (for DHCPv6) to the DHCP client. The relay agent cannot tear down the binding itself, so it sends a release message to the local server, causing the local server to tear down the binding. For this to happen, you must configure the `send-release-on-delete` statement on the DHCP relay agent; otherwise it will not send the release message to the local server. In that case, the local server retains the client entry in the database until it times out or the IP address is used for a different binding.

NOTE: `remote-id-mismatch` functionality overrides the default DHCP relay agent bind-on-request behavior. By default, when a stray DHCP request is received, that is, one for which there is an entry in the local server database but not in the relay agent database, a complete binding is automatically made with the relay agent and the local server.

The DHCP client initiates renegotiation when it receives the NAK. The changed Agent Remote ID value is conveyed as part of the request, enabling the new service plan to be submitted for authorization.

The `remote-id-mismatch` statement is typically used in an environment that uses local authorization instead of RADIUS authorization.

NOTE: 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.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring DHCP-Initiated Service Change Based on Remote ID | 745

replace-ip-source-with (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2306
- Hierarchy Level | 2306
- Description | 2306
- Required Privilege Level | 2306
- Release Information | 2306

Syntax

```
replace-ip-source-with giaddr;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address | 717

report-interface-descriptions (Access)

IN THIS SECTION

- Syntax | 2307
- Hierarchy Level | 2307
- Description | 2307
- Required Privilege Level | 2308
- Release Information | 2308

Syntax

```
report-interface-descriptions;
```

Hierarchy Level

```
[edit access]
```

Description

Enable storing and reporting of interface descriptions through RADIUS. To disable storing and reporting of interface descriptions, configure the `[edit access profile` *`profile-name`* `radius attributes exclude]` statement to exclude the interface description attribute. The description can contain letters, numbers, and hyphens (-), and can be up to 64 characters long.

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 13.1.

RELATED DOCUMENTATION

- Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers | 498
- RADIUS Servers and Parameters for Subscriber Access | 476
- Configuring Authentication and Accounting Parameters for Subscriber Access | 550

report-local-rule (PCRF Partition)

IN THIS SECTION

- Syntax | 2308
- Hierarchy Level | 2309
- Description | 2309
- Required Privilege Level | 2309
- Release Information | 2309

Syntax

```
report-local-rule;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to send local report downstream messages by default. If you configure this statement, this generates the same result as if you enabled the Resource-Allocation-Notification message trigger with the local rule definition.

The following example shows predefined rule installation commands, including the Resource-Allocation-Notification command, for Credit Control Answer (CCA-GX) and Reauthorization Request (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) ]
 }
```

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | 1455

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
Understanding Gx Interactions Between the Router and the PCRF | 1418
Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
Understanding Upstream and Downstream Messages for the PCRF | 1444

report-resource-allocation (PCRF Partition)

IN THIS SECTION

- Syntax | 2310
- Hierarchy Level | 2310
- Description | 2310
- Required Privilege Level | 2311
- Release Information | 2311

Syntax

```
report-resource-allocation;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to generate reports for resource allocation failures by default. If you configure this statement, this generates the same result as if you enabled the Resource-Allocation-Notification message trigger with every rule received from the PCRF.

The following example shows predefined rule installation commands, including the Resource-Allocation-Notification command, for Credit Control Answer (CCA-GX) and Reauthorization Request (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 provides generated reports by default.

For login and update interactions between the PCRF and the Online Charging System (OCS), the OCS sends applicable reports to the PCRF using a CCR-GX-U message. If the service-dynamic-profile instantiation fails, and the Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule, then the OCS sends a report. The router acknowledges the subscriber activation, allows subscriber traffic to flow, and continues to repeat the process.

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | 1455

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
Understanding Gx Interactions Between the Router and the PCRF | 1418
Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
Understanding Upstream and Downstream Messages for the PCRF | 1444

report-successful-resource-allocation (PCRF Partition)

IN THIS SECTION

- Syntax | 2312
- Hierarchy Level | 2312
- Description | 2312
- Required Privilege Level | 2313
- Release Information | 2313

Syntax

```
report-successful-resource-allocation;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to generate reports for both resource allocation successes and failures by default. If you configure this statement, this generates the same result as if you enabled the Successful-Resource-Allocation-Notification message trigger with every CCA-GX-I, CCA-GX-U, and RAR-GX-A message.

The following example 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 Policy and Charging Enforcement Function (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 provides generated reports by default.

For login and update interactions between the PCRF and the Online Charging System (OCS), the OCS sends applicable reports to the PCRF using a CCR-GX-U message. If the service-dynamic-profile instantiation succeeds, and the Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule, and the SUCCESSFUL_RESOURCE_ALLOCATION event trigger is set in the request, then the OCS sends a report.

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 16.2.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | **1455**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**
- Understanding Upstream and Downstream Messages for the PCRF | **1444**

request-max-tcp-connections (System Process)

IN THIS SECTION

- Syntax | 2314
- Hierarchy Level | 2314
- Description | 2314
- Options | 2314
- Required Privilege Level | 2315
- Release Information | 2315

Syntax

```
request-max-tcp-connections max-tcp-connections;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Specify the number of simultaneous TCP connections that the DHCP relay can request when sending bulk leasequery messages to DHCP servers.

Options

max-tcp-connections

Number of simultaneous connections.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

system—to view this statement in the configuration.

system-control—to add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery | 812

Configuring and Using DHCP Bulk Leasequery | 814

request-rate (Access)

IN THIS SECTION

- Syntax | 2315
- Hierarchy Level | 2316
- Description | 2316
- Options | 2316
- Required Privilege Level | 2316
- Release Information | 2316

Syntax

```
request-rate rate;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the number of requests the router can send per second to all configured RADIUS servers collectively. By limiting the flow of requests from the router to the RADIUS servers, you can prevent the RADIUS servers from being flooded with requests.

Options

rate—Number of requests per second.

- **Range:** 100 through 4000 requests per second
- **Default:** 500 requests per second

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 11.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

requested-ip-network-match (DHCP Local Server)

IN THIS SECTION

- Syntax | 2317
- Hierarchy Level | 2317
- Description | 2317
- Options | 2318
- Required Privilege Level | 2318
- Release Information | 2318

Syntax

```
requested-ip-network-match subnet-mask;
```

Hierarchy Level

```
[edit system services dhcp-local-server]
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure the subnet to which the DHCP local server matches the requested IP address (IPv6 address for DHCPv6 local server). The server accepts and uses the active client's requested address for address assignment only when the requested address and the IP address of the DHCP server interface (or IPv6 address of the DHCPv6 local server) are in the same subnet. The server accepts and uses the passive client's requested address for address assignment only when the requested address and the address of the relay interface are in the same subnet.

Options

subnet-mask	Length of the subnet mask.
	• Range:
	  • DHCP: 0 through 31
	  • DHCPv6: 0 through 127
	• Default:
	  • DHCP: 8
	  • DHCPv6: 16

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Specifying the Subnet for DHCP Client Address Assignment | 776

retransmit-timer (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2319
- Hierarchy Level | 2319
- Description | 2319
- Options | 2319

Required Privilege Level | 2319
Release Information | 2319

Syntax

```
retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Set the retransmission frequency of neighbor solicitation messages.

Options

milliseconds—Retransmission frequency.

- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

revert-interval (Access)

IN THIS SECTION

- Syntax | 2320
- Hierarchy Level | 2320
- Description | 2320
- Options | 2320
- Required Privilege Level | 2321
- Release Information | 2321

Syntax

```
revert-interval interval;
```

Hierarchy Level

```
[edit access profile profile-name radius options],
[edit access radius-options]
```

Description

Configure the amount of time the router or switch waits after a server has become unreachable. The router or switch rechecks the connection to the server when the specified interval expires. If the server is then reachable, it is used in accordance with the order of the server list.

Options

interval—Amount of time to wait.

- **Range:** 0 through 604,800 seconds
- **Default:** 60 seconds

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 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

RADIUS Authentication and Accounting Basic Configuration | 550

route (Diameter Network Element)

IN THIS SECTION

- Syntax | 2321
- Hierarchy Level | 2322
- Description | 2322
- Options | 2322
- Required Privilege Level | 2322
- Release Information | 2322

Syntax

```
route dne-route-name {
    destination realm realm-name  <host hostname>;
```

```
    function function-name <partition partition-name>;
    metric route-metric;
}
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding]
```

Description

Define a route reachable through the Diameter network element by associating a metric with a combination of destination and function partition.

Options

dne-route-name—Route name defined for the Diameter network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Network Elements | 1377

router-advertisement (Dynamic Profiles)

IN THIS SECTION

- Syntax | 2323
- Hierarchy Level | 2323
- Description | 2324
- Required Privilege Level | 2324
- Release Information | 2324

Syntax

```
router-advertisement {
    interface interface-name {
         current-hop-limit number;
        default-lifetime seconds;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefixprefix {
            (autonomous | no-autonomous);
            (on-link | no-on-link);
            preferred-lifetime seconds;
            valid-lifetime seconds;
        }
        reachable-time milliseconds;
        retransmit-timer milliseconds;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols]
```

Description

Enable router advertisement in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

routing-instance (Diameter Peer)

IN THIS SECTION

- Syntax | 2325
- Hierarchy Level | 2325
- Description | 2325
- Options | 2325
- Required Privilege Level | 2325
- Release Information | 2325

Syntax

```
routing-instance routing-instance-name ;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify a routing instance for a Diameter peer. Alternatively, you can include the `logical-system` statement at the `[edit diameter peer peer-name]` hierarchy level to configure a logical and routing instance.

Options

routing-instance-name—Name of the routing instance.

- **Default:** Master routing instance

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 9.6.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring Diameter Peers | 1374

routing-instance (Diameter Transport)

IN THIS SECTION

- Syntax | 2326
- Hierarchy Level | 2326
- Description | 2326
- Options | 2326
- Required Privilege Level | 2326
- Release Information | 2327

Syntax

```
routing-instance routing-instance-name ;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Specify a routing instance for the Diameter transport layer connection.

Options

routing-instance-name—Name of the routing instance.

- **Default:** Master routing instance

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 11.2.

RELATED DOCUMENTATION

Configuring the Diameter Transport | 1376

routing-instance-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 2327
- Hierarchy Level | 2327
- Description | 2328
- Required Privilege Level | 2328
- Release Information | 2329

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the routing instance name be concatenated with the username during the subscriber authentication or DHCP client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

routing-instance-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2329
- Hierarchy Level | 2329
- Description | 2330
- Required Privilege Level | 2330
- Release Information | 2330

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit tenants tenant-name routing-instances routing-instance-name forwarding-options dhcp-relay
dhcpv6 group group-name authentication username-include]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify that the routing instance name is concatenated with the username during the subscriber authentication or client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.

Support at the [edit ... dual-stack-group *dual-stack-group-name*] hierarchy level introduced in Junos OS Release 15.1.

The tenants option is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

routing-instance-name (Static Subscribers)

IN THIS SECTION

- Syntax | 2331
- Hierarchy Level | 2331
- Description | 2332
- Required Privilege Level | 2332
- Release Information | 2332

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
```

```
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that the name of the routing instance is included as part of the username created for all static subscribers or for the static subscribers in the specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487
- Configuring the Static Subscriber Global Username | 1492
- Configuring the Static Subscriber Group Username | 1497

s6a

IN THIS SECTION

- Syntax | 2333
- Hierarchy Level | 2333
- Description | 2333
- Options | 2334
- Required Privilege Level | 2334
- Release Information | 2334

Syntax

```
s6a {
    partition name {
        destination-host destination-host;
        destination-realm destination-realm;
        diameter-instance diameter-instance;
        max-outstanding-requests max-outstanding-requests;
        response-timeout seconds;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

S6a is a Diameter-based authentication application used by the Mobility Management Entity (MME) nodes to retrieve authentication information from Home Subscriber Server (HSS). Specify the destination and transmission parameters for the S6a protocol.

Options

destination-host Name of the host where the S6A server application resides; generally, this value is not set unless needed.

destination-realm Name of the realm where the S6a server application resides.

diameter-instance Specifies that the master Diameter instance is used. `master` is the only supported value.

max-outstanding-requests Configure the maximum number of outstanding requests for the S6a application.

- **Default:** 40
- **Range:** 2 through 1024

response-timeout Configure the amount of time (in seconds) that the MME waits to receive a response from the HSS.

- **Default:** 15
- **Range:** 1 through 30

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configure S6a Application | **1379**

authentication-order | **1640**

partition (s6a) | **2166**

Diameter Base Protocol Overview | **1339**

sdsl-bytes

IN THIS SECTION

- Syntax | 2335
- Hierarchy Level | 2335
- Description | 2335
- Options | 2336
- Required Privilege Level | 2336
- Release Information | 2336

Syntax

```
sdsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an SDSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `sdsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

sdsl-overhead-adjust

IN THIS SECTION

Syntax | 2337

- Hierarchy Level | 2337
- Description | 2337
- Options | 2337
- Required Privilege Level | 2338
- Release Information | 2338

Syntax

```
sdsl-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for an SDSL access line reported in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `sdsl-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299
- Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293
- Configuring the ANCP Agent | 1253

send-acct-status-on-config-change (Access Profile)

IN THIS SECTION

- Syntax | 2339
- Hierarchy Level | 2339
- Description | 2339
- Required Privilege Level | 2339
- Release Information | 2339

Syntax

```
send-acct-status-on-config-change;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router's authd process to send accounting messages when the RADIUS server status changes for an access profile. When you include this statement, authd sends an Acct-On message when the first RADIUS server is added to an access profile, and to send an Acct-Off message when the last RADIUS server is removed from an access 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 13.1.

RELATED DOCUMENTATION

Configuring Per-Subscriber Session Accounting | 573

RADIUS Servers and Parameters for Subscriber Access | 476

send-dyn-subscription-indicator (PCRF Partition)

IN THIS SECTION

- Syntax | 2340
- Hierarchy Level | 2340
- Description | 2340
- Required Privilege Level | 2341
- Release Information | 2341

Syntax

```
send-dyn-subscription-indicator;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
```

Description

Include the Juniper-Network-Family-Indicator AVP into its message requests. To enable the Policy Control and Charging Rules Function (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 92 on page 2341 have been allocated from the Juniper Vendor-ID space (2636) vendor-specific attribute (VSA).

NOTE: If no dynamic-subscription ID is received, then neither Online Charging System (OCS) or Offline Charging System (OFCS) communications is initiated.

Table 92: 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.

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | **1455**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | 1418
Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
Understanding Upstream and Downstream Messages for the PCRF | 1444

send-network-family-indicator (PCRF Partition)

IN THIS SECTION

- Syntax | 2342
- Hierarchy Level | 2342
- Description | 2342
- Required Privilege Level | 2343
- Release Information | 2343

Syntax

```
send-network-family-indicator;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Include the Juniper-Network-Family-Indicator AVP into its message requests. In many cases, wireline subscribers support only one IP family, which is required information for both AAA service and Policy Control and Charging Rules Function (PCRF). To indicate this information, the Juniper-Network-Family-Indicator AVP has been allocated from the Juniper Vendor-ID space (2636) VSA in Table 93 on page 2343.

Table 93: Network 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)

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | **1455**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | **1418**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

Understanding Upstream and Downstream Messages for the PCRF | **1444**

send-origin-state-id (Diameter Applications)

IN THIS SECTION

- Syntax | 2344
- Hierarchy Level | 2344
- Description | 2344
- Required Privilege Level | 2345
- Release Information | 2345

Syntax

```
send-origin-state-id;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name],
[edit diameter peer peer-name]
```

Description

Include the Origin-State AVP for the Diameter peer, OCS partition, or PCRF partition. The AVP is conveyed in Diameter-level or partition-level messages, respectively. The value of the AVP increases every time it changes, such as when the system reboots. Sending the Origin-State AVP in messages enables Diameter entities that receive the message to infer from a changed AVP value that sessions associated with a lower value are no longer active.

If you include the `send-origin-state-id` statement for the OCS or PCRF partition, a new value is sent for the partition only when a cold boot occurs. The system daemon detects the cold boot and changes the Origin-State-Id when required.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | 1449
- Configuring the PCRF Partition | 1455
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

send-release-on-delete (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2345
- Hierarchy Level | 2346
- Description | 2346
- Required Privilege Level | 2346
- Release Information | 2347

Syntax

```
send-release-on-delete;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Send a release message to the DHCP (or DHCPv6) server whenever DHCP relay or relay proxy deletes a client. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Extended DHCP Relay Agent Overview | 695
- Overriding the Default DHCP Relay Configuration Settings | 709
- Sending Release Messages When Clients Are Deleted | 713

server-duid-type (DHCP Local Server)

IN THIS SECTION

- Syntax | 2347
- Hierarchy Level | 2347
- Description | 2348
- Options | 2348
- Required Privilege Level | 2348
- Release Information | 2348

Syntax

```
server-duid-type type;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6]
```

Description

Configure the DHCP Unique Identifier (DUID) type that the DHCPv6 server supports. The server uses the DUID to identify clients for configuration and to group clients into identity associations.

Use this statement to explicitly configure the DUID-LL type. DUID-LL uses a link-layer network interface address on the device as part of the unique identifier. Remove this configuration to return to the default type, DUID-EN. DUID-EN has a vendor-assigned value based on the vendor's enterprise number and an IANA-recognized hardware type code. DUID-EN is not explicitly configurable.

NOTE: The DUID-LLT type is not supported.

Options

type DUID type to support. The only available option is `duid-ll`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCPv6 Local Server Overview | 910

server-group

IN THIS SECTION

Syntax | 2349

- Hierarchy Level | 2349
- Description | 2349
- Options | 2350
- Required Privilege Level | 2350
- Release Information | 2350

Syntax

```
server-group {
    server-group-name {
        server-ip-address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6]
```

Description

Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Apply the group with the `active-server-group` statement globally for all interfaces or for a named group of interfaces configured with the `group` statement. This mechanism enables you to apply different DHCP relay configurations for different groups of servers, with a common configuration for the servers within a server group.

Options

server-group-name	Name of the group of DHCP or DHCPv6 server addresses.
server-ip-address	IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. Starting in Junos OS Release 18.4R1, you can configure up to 32 server IP addresses per group for DHCPv4 servers. In earlier releases, you can configure only up to 5 server IP addresses for DHCPv4 servers. For DHCPv6 servers, you can configure only up to 32 addresses in all releases. The configuration fails commit check if you configure more than the maximum number of server addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview | 695

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups | 858

server-id-override

IN THIS SECTION

Syntax | 2351

Hierarchy Level | 2351

- Description | 2351
- Required Privilege Level | 2352
- Release Information | 2352

Syntax

```
server-id-override;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-option-82]
```

Description

Add the suboptions `link-selection` and `server-id override` to option-82 information for DHCP packets relayed to the server.

In network configurations where a firewall on the broadband network gateway (BNG) is between the DHCPv4 relay agent and the DHCP server, only the BNG loopback address passes through the BNG firewall. In that case, DHCP unicast packets do not pass through the firewall and and discarded. To enable DHCP unicast packets to pass through the BNG firewall, use the `relay-source` configuration statement to configure the DHCP relay agent to use the loopback address as the source address in IP headers and DHCP messages. In this case, this configuration statement adds the `link-selection` suboption to provide the DHCP server information to locate the correct address pool for the DHCP client and adds the `server-id override` suboption to set the server ID option.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall | 717

server-response-time (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2352
- Hierarchy Level | 2353
- Description | 2353
- Options | 2353
- Required Privilege Level | 2353
- Release Information | 2353

Syntax

```
server-response-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Configure the timeframe during which the router monitors DHCP server responsiveness within the routing instance. The router generates a system log message when the DHCP server does not respond to relayed packets during the specified timeframe.

Options

seconds Number of seconds the DHCP server is monitored.

- **Range:** 30 through 4,294,967,295 seconds
- **Default:** 0 (no limit)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Monitoring DHCP Relay Server Responsiveness | 898

service (Service Accounting)

IN THIS SECTION

- Syntax | 2354
- Hierarchy Level | 2354
- Description | 2354
- Required Privilege Level | 2354
- Release Information | 2355

Syntax

```
service {
    accounting-order (activation-protocol | local | radius);
    accounting {
        statistics (time | volume-time);
        update-interval minutes;
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Define the subscriber service accounting configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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 11.4.

`accounting`, `update-interval`, and `statistics` options added in Junos OS Release 14.2R1 for MX Series routers.

RELATED DOCUMENTATION

- Configuring Service Accounting with JSRC | **1482**
- Service Accounting with JSRC | **1480**
- *Configuring Service Accounting in Local Flat Files*
- Configuring Service Accounting | **586**
- Configuring Per-Subscriber Session Accounting | **573**

service-context-id (OCS)

IN THIS SECTION

- Syntax | **2355**
- Hierarchy Level | **2356**
- Description | **2356**
- Options | **2356**
- Required Privilege Level | **2356**
- Release Information | **2357**

Syntax

```
service-context-id service-context ;
```

Hierarchy Level

```
[edit access ocs global]
```

Description

Configure the service-context-id to globally and uniquely identify the Service-Context-Id AVP (cc-service-context) and the Service-Identifier AVP (cc-service-identifier), which is part of the Diameter Credit Control Service charging system. The service provider or operator allocates this identifier.

Options

service-context 3rd Generation Partnership Project (3GPP)-specific values.

- *service-context*—32251
- *customer-domain*—3gpp.org

> **NOTE**: For every service used to support the packet-switching charging infrastructure, you must prepare the necessary documentation and define associated cc-service-context-id values.

For mobile subscribers, the user equipment requests services; whereas for broadband wireline subscribers, the Policy Control and Charging Rules Function (PCRF) requests services. In the wireline environment, PCRF functions as the service requester, and the Policy and Charging Enforcement Function (PCEF) functions as the service receiver and enforcer.

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, you must configure the PCEF to define a charging rules mapping between credit control services (cc-services) and rating groups.

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 16.2.

RELATED DOCUMENTATION

- Configuring OCS Global Parameters | 1462
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | 1410
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

service-profile (DHCP Local Server)

IN THIS SECTION

- Syntax | 2357
- Hierarchy Level | 2357
- Description | 2358
- Options | 2358
- Required Privilege Level | 2358
- Release Information | 2358

Syntax

```
service-profile dynamic-profile-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
```

```
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[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 ...]
```

Description

Specify the default subscriber service or DHCP client management service, which is activated when the subscriber or client logs in and no other service is activated by a RADIUS server or a provisioning server.

- To specify the default service for all DHCP local server clients, include the `service-profile` statement at the `[edit system services dhcp-local-server]` hierarchy level.
- To specify the default service for a named group of interfaces, include the `service-profile` statement at the `[edit system services dhcp-local-server group group-name]` hierarchy level.
- To specify the default service for a particular interface within a named group of interfaces, include the `service-profile` statement at the `[edit system services dhcp-local-server group group-name interface interface-name]` hierarchy level.
- For DHCPv6 clients, use the `service-profile` statement at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

Options

dynamic-profile-name—Name of the dynamic profile that defines the service.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP | 691

Default Subscriber Service Overview | 764

Configuring a Default Subscriber Service | 765

service-profile (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2359
- Hierarchy Level | 2359
- Description | 2360
- Options | 2360
- Required Privilege Level | 2360
- Release Information | 2360

Syntax

```
service-profile dynamic-profile-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit forwarding-options dhcp-relay dhcpv6 groupgroup-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the default subscriber service (or the default DHCP client management service), which is activated when the subscriber (or client) logs in and no other service is activated by a RADIUS server or a provisioning server.

- To specify the default service for all DHCP relay agent clients, include the `service-profile` statement at the `[edit forwarding-options dhcp relay]` hierarchy level.
- To specify the default service for a named group of interfaces, include the `service-profile` statement at the `[edit forwarding-options dhcp relay group group-name]` hierarchy level.
- To specify the default service for a particular interface within a named group of interfaces, include the `service-profile` statement at the `[edit forwarding-options dhcp relay group group-name interface interface-name]` hierarchy level.

Options

dynamic-profile-name—Name of the dynamic service profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 1752

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Grouping Interfaces with Common DHCP Configurations | 852

Default Subscriber Service Overview | 764

service-profile (Static Subscribers)

IN THIS SECTION

- Syntax | 2361
- Hierarchy Level | 2361
- Description | 2361
- Options | 2361
- Required Privilege Level | 2362
- Release Information | 2362

Syntax

```
service-profile service-profile-name;
```

Hierarchy Level

```
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
```

Description

Specify the service profile to apply services for all static subscribers at the global level and at the group level. You can use `service-profile` to assign a default service profile to be applied to a `static-subscriber` session if the policy server is unavailable.

Options

service-profile-name—Name of the service profile.

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 17.4R1.

RELATED DOCUMENTATION

- Subscribers over Static Interfaces Configuration Overview | 1487
- Specifying the Static Subscriber Group Service Profile | 1495
- Specifying the Static Subscriber Global Access Profile | 1490

services (System Services)

IN THIS SECTION

- Syntax | 2362
- Hierarchy Level | 2369
- Description | 2369
- Required Privilege Level | 2369
- Release Information | 2369

Syntax

```
services {
    dhcp { # DHCP is not supported on a DCF
        dhcp_services;
    }
    dtcp-only
    finger {
```

```
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
    ftp {
        authentication-order [authentication-methods];
        connection-limit limit;
        rate-limit limit;
    }
    grpc {
        request-response {
            grpc {
                ssl {
                    address ip-address;
                    local-certificate local-certificate;
                    port port;
                }
                max-connections max-connections;
            }
        }
        notification {
            port port;
            max-connections max-connections;
            allow-clients {
                address ip-address;
            }
        }
        traceoptions {
            file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
            flag flag;
            no-remote-trace;
        }
    }
    netconf {
        flatten-commit-results;
        hello-message {
            yang-module-capabilities {
```

```
                advertise-native-yang-modules;
                advertise-custom-yang-modules;
                advertise-standard-yang-modules;
            }
        }
        netconf-monitoring {
            netconf-state-schemas {
                retrieve-custom-yang-modules;
                retrieve-standard-yang-modules;
            }
        }
        notification;
        rfc-compliant;
        ssh {
            client-alive-count-max number;
            client-alive-interval seconds;
            connection-limit limit;
            port port;
            rate-limit limit;
        }
        tls {
            client-identity client-id {
                fingerprint fingerprint;
                map-type (san-dirname-cn | specified);
                username username;
            }
            default-client-identity {
                map-type (san-dirname-cn | specified);
                username username;
            }
            local-certificate local-certificate;
            traceoptions {
                file <filename> <files files> <match match> <size size> <(world-readable | no-
world-readable)>;
                flag name;
                level (all | error | info | notice | verbose | warning);
                no-remote-trace;
            }
        }
        traceoptions {
            file <filename> <files number> <match regular-expression> <size size> <world-
readable | no-world-readable>;
            flag flag;
```

```
            no-remote-trace;
            on-demand;
        }
        yang-compliant;
        yang-modules {
            device-specific;
            emit-extensions;
        }
    }
    outbound-https {
        client client-id {
            address {
                port port;
                trusted-cert trusted-cert;
            }
            device-id device-id;
            reconnect-strategy (in-order | sticky);
            secret password;
            waittime seconds;
        }
    }
    service-deployment {
        servers address {
            port-number port-number;
        }
        source-address address;
    }
    ssh {
        authentication-order  [method 1 method2...];
        authorized-keys-command authorized-keys-command;
        authorized-keys-command-user authorized-keys-command-user;
        ciphers [ cipher-1 cipher-2 cipher-3 ...];
        client-alive-count-max number;
        client-alive-interval seconds;
        connection-limit limit;
        fingerprint-hash (md5 | sha2-256);
        hostkey-algorithm (algorithm / no-algorithm);
        key-exchange [algorithm1 algorithm2...];
        log-key-changes log-key-changes;
        macs [algorithm1 algorithm2...];
        max-pre-authentication-packets number;
        max-sessions-per-connection number;
        no-challenge-response;
```

```
        no-password-authentication;
        no-passwords;
        no-public-keys;
        allow-tcp-forwarding;
        port port-number;
        protocol-version [v2];
        rate-limit number;
        rekey {
            data-limit bytes;
            time-limit minutes;
        }
        root-login (allow | deny | deny-password);
        sftp-server;
        }
    tcp-forwarding;
    resource-monitor {
        free-fw-memory-watermark number;
        free-heap-memory-watermark number;
        free-nh-memory-watermark number;
        high-threshold number;
        no-logging;
        no-throttle;
        resource-category jtree {
            resource-category jtree (continguous-pages | free-dwords | free-pages) {
                low-watermark number;
                high-watermark number;
            }
        }
        subscribers-limit {
             (any | dhcp | l2tp | pppoe) {
                 {
                    limit limit;
                }
{
                    limit limit;
                }
                fpc slot-number {
                    limit limit;
                    pic number {
                        limit limit;
                        port number {
                            limit limit;
                        }
```

```
                    }
                }
            }
        }
        traceoptions {
            file filename <files number> <match regular-expression> <size maximum-file-size>
<world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    subscriber-management {
        enable (Enhanced Subscriber Management);
        enforce-strict-scale-limit-license;
        gres-route-flush-delay;
        }
        overrides {
            event {
                catastrophic-failure {
                    reboot (master | standby);
                }
            }
            interfaces {
                family (inet | inet6) {
                    layer2-liveness-detection;
                }
            }
            no-unsolicited-ra;
            ra-initial-interval-max seconds;
            ra-initial-interval-min seconds;
            shmlog {
                disable;
                file filename <files maximum-no-files> <size maximum-file-size>;
                filtering enable;
                log-name {
                    all;
                    logname {
                        <brief | detail | extensive | none | terse>;
                        <file-logging |no-file-logging>;
                    }
                    }
                log-type (debug | info | notice);
            |
```

```
        }
        redundancy {
            interface name {
                local-inet-address v4-address;
                local-inet6-address v6-address;
                shared-key string;
                virtual-inet-address virtual-v4-address;
                virtual-inet6-address virtual-v6-address;
            }
            no-advertise-routes-on-backup;
            protocol {
                pseudo-wire;
                vrrp;
            }
        }
        traceoptions {
            file filename <files number> <match regular-expression > <size maximum-file-size>
<world-readable | no-world-readable>;
            flag flag;
        }
    }
    telnet {
        authentication-order [authentication-methods];
        connection-limit limit;
        rate-limit limit;
    }
    web-management {
        http {
            interfaces [ names ];
            port port;
        }
        https {
            interfaces [ names ];
            local-certificate name;
            port port;
        }
        session {
            idle-timeout [ minutes ];
            session-limit [ limit ];
        }
    }
    xnm-ssl {
        connection-limit limit;
```

```
        local-certificate name;
        rate-limit limit;
        ssl-renegotiation;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the router or switch so that users on remote systems can access the local router or switch through the DHCP server, DTCP over SSH, finger, outbound HTTPS, rlogin, SSH, telnet, Web management, Junos XML protocol SSL, and network utilities, or enable Junos OS to work with the Session and Resource Control (SRC) software. Also, enable configuration of third-party applications developed using the Juniper Extension Toolkit (JET) to run on Junos OS.

Starting in Junos OS Release 22.2R1, we've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the [`edit system services ssh`] hierarchy level. In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the [`edit system services ssh`] hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`extension-service` option added in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.

`grpc` option added in Junos OS Release 16.2 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.

`allow-tcp-forwarding` option added in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

- Configuring the Junos OS to Work with SRC Software
- M:N Subscriber Redundancy on BNG Overview | **1173**

session-limit-per-username (Access Profile)

IN THIS SECTION

- Syntax | **2370**
- Hierarchy Level | **2370**
- Description | **2370**
- Options | **2371**
- Required Privilege Level | **2371**
- Release Information | **2371**

Syntax

```
session-limit-per-username number;
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Limit the number of active sessions allowed for a username associated with a given access profile. This enables you to control credential sharing by legitimate subscribers. After the limit is configured, existing

subscriber sessions for the access profile are tracked as active sessions, as are any new sessions that do not exceed the limit. Requests for sessions that exceed the configured limit are blocked.

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.

Options

number Maximum number of subscriber sessions allowed per username in the access profile.

- **Range:** 1 through 16

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access | 504

Limiting the Number of Active Sessions per Username and Access Profile | 512

session-options

IN THIS SECTION

- Syntax | 2372
- Hierarchy Level | 2373
- Description | 2373
- Options | 2373
- Required Privilege Level | 2375
- Release Information | 2375

Syntax

```
session-options {
    client-group [ group-names ];
    client-idle-timeout minutes;
    client-idle-timeout-ingress-only;
    client-session-timeoutminutes;
    pcc-context {
        input-service-filter-name filter-name;
        input-service-set-name service-set-name;
        ipv6-input-service-filter-name filter-name;
        ipv6-input-service-set-name service-set-name;
        ipv6-output-service-filter-name filter-name;
        ipv6-output-service-set-name service-set-name;
        output-service-filter-name filter-name;
        output-service-set-name service-set-name;
        profile-name pcef-profile-name;
    }
    strip-user-name {
        delimiter [ delimiter ];
        parse-direction (left-to-right | right-to-left);
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

(MX Series and SRX Series devices) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

(MX Series) Specify characteristics related to policy and charging control (PCC) rules, such as the PCEF profile that contains the rules, service sets to process the rules, and service filters for the service sets.

Options

client-idle-timeout Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

During this period, the router determines whether the subscriber is inactive by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, non-DHCP subscribers (such as L2TP or PPP) are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout; DHCP subscribers are disconnected.

When you additionally configure the related `client-idle-timeout-ingress-only` statement (MX Series only), the router monitors only ingress traffic to determine whether the subscriber is inactive; it does not monitor any egress traffic. The related `client-session-timeout` statement terminates the subscriber session when the session timeout expires regardless of user activity.

Client idle timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model. It is not practical for DHCP or DHCPv6 subscribers.

Although you can use the `client-idle-timeout` statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the idle timeout for VLANs, the timeout period starts when the VLAN is instantiated. It

resets when a client session is created or an existing session is reactivated. When no traffic is detected on an authenticated VLAN for the duration of the timeout, the VLAN is considered inactive and is deleted. If no client sessions are ever created on the VLAN, then the VLAN is removed when the timeout expires.

- **Default:** The timeout is not configured.
- **Values:** *minutes*—Number of minutes of idle time that elapse before the session is terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.
- **Range:** 10 through 1440 minutes

client-idle-timeout-ingress-only

Specify that only ingress traffic is monitored for subscriber idle timeout processing for the duration of the idle timeout period that you specify with the `client-idle-timeout` statement. If no ingress traffic is received for the duration of the timeout, then the subscriber is gracefully logged out (non-DHCP subscribers) or disconnected (DHCP subscribers).

If you configure `client-idle-timeout` alone, then both ingress and egress traffic are monitored during the idle timeout. Monitoring only ingress traffic is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore does not detect that the peer is not up. Because the LAC monitors both ingress and egress traffic by default, in this situation it receives the egress traffic from the LNS and either does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored in this case, the LAC can detect that the peer is inactive and then initiate logout.

client-session-timeout

Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Alternatively, when you want subscribers to be identified as inactive before they are terminated, use the related statements, `client-idle-timeout` and `client-idle-timeout-ingress-only`. Use `client-idle-timeout` alone to specify a period of time during which both ingress and egress subscriber data traffic is monitored; if no traffic is detected for the duration of the period, the subscriber is considered inactive and is terminated. Add the `client-idle-timeout-ingress-only` statement to monitor only ingress traffic for the duration of the timeout set with the `client-idle-timeout statement`.

BEST PRACTICE: We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is a simple time-based timeout, it is likely to interrupt subscribers actively using a voice

service and terminate their calls unexpectedly (from the subscriber viewpoint). This result is a particular concern for emergency services calls.

Client session timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model when no voice services are offered. For DHCP or DHCPv6 subscribers, the session timeout is used as the DHCP lease timer if no other lease time configuration is present.

Although you can use the `client-session-timeout` statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the session timeout for VLANs, the timeout period starts when the VLAN is instantiated.

- **Default:** The timeout is not configured.
- **Values:** *minutes*—Number of minutes after which user sessions are terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.
- **Range:** 1 through 527040 minutes

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Session Options for Subscriber Access | 503

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

sftp-backup (OCS Partition)

IN THIS SECTION

- Syntax | 2376
- Hierarchy Level | 2377
- Description | 2377
- Options | 2377
- Required Privilege Level | 2378
- Release Information | 2378

Syntax

```
sftp-backup {
    accumulation-timeout seconds;
    accumulation-count messages;
    accumulation-size bytes;
    logical-system;
    retry-interval seconds;
    response-timeout seconds;
    routing-instance;
    address ipv4 | ipv6;
    port number;
    directory filename;
    file-name-prefix filename-prefix;
    node-ipv4-address;
    node-ipv6-address;
    ssh-login;
    ssh-connection-linger;
    ssh-log-verbose;
    ssh-passphrase;
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

The data backup for OCS using the SFTP protocol to store data on a remote server.

Options

accumulation-timeout *seconds*	The file accumulation time since the first CCR-GY-T was submitted. The maximum time is 30 seconds.
accumulation-count *messages*	The number of requests that are past file account count. The maximum accumulation count is 2000 messages.
accumulation-size *bytes*	The file once its size has reached the accumulation size limit. The maximum accumulation size is 3000000 bytes.
logical-system	The logical system to be used by sftp.
retry-interval *seconds*	Every failed write operation is retried after this interval until backup timeout is accumulated. The maximum retry interval is 1200 seconds.
response-timeout *seconds*	The timeout on individual SFTP command response. The maximum response timeout is 60 seconds.
routing-instance	The SFTP service connection context for the routing instance.
address	The address of SFTP server
port	The port number. The default port number is 22.
directory	The directory name.
file-prefix *prefix*	The prefix for the file.
node-ipv4-address	The IPv4 address of the sending node.
node-ipv6-address	The IPv6 address of the sending node.

ssh-login	The login name on SFTP server, may include %h and/or %r similar to directory.
ssh-connection-linger	The time to keep existing ssh connection idle state.
ssh-log-verbose	The verbose logs to be collected from ssh session.
ssh-passphrase	The password or later passphrase.

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 18.1.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- Gy File Backup Overview | **1438**
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

shmlog (Shared Memory Log)

IN THIS SECTION

- Syntax | **2379**
- Hierarchy Level | **2379**
- Description | **2379**

- Options | 2380
- Required Privilege Level | 2381
- Release Information | 2381

Syntax

```
shmlog {
    disable;
    file filename <files maximum-no-files> <size maximum-file-size>;
    filtering enable;
    log-name {
        all;
        logname {
            <brief | detail | extensive | none | terse>;
            <file-logging |no-file-logging>;
        }
        }
    log-type (debug | info | notice);
|
```

Hierarchy Level

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

Description

Junos OS uses a shared memory space to store log entries for subscriber service daemons including jpppd, jdhcpd, jl2tpd, autoconfd, bbe-smgd, authd, cosd, and dfwd. Shared memory logging is enabled by default and occurs at the client level. You can view the shmlogs on a per subscriber basis, or use filters to retrieve logs according to a variety of different parameters such as interface name, IP address, session ID, subnet, and VLAN in addition to the Client Identifier or Client DUID. Filtering is disabled by default. To see a complete list of supported filters, use this command:

```
user@rdevice> show shmlog entries logname all ?
```

When viewing logs you can limit results on the basis of event flags that include interface events, routing process interaction events, l2tp tunneling events, and ldap authentication events. To see a complete list of supported flags, use this command:

```
user@rdevice> show shmlog entries logname all flag-name ?
```

NOTE: Some platforms other than MX Series routers use shared memory logs for internal processes. These logs are not intended for customer use.

Options

disable Name of the command to override the default behavior. Use this option to disable shared memory logging; it is always enabled otherwise.

file Name of the file containing the shmlogs. Use this option to redirect shmlogs to a file for file-based logging. Specify the file name, define the number of files (from 2 to 1000), and set the maximum file size (from 10240 to 1073741824 bytes). Data will be written to the **/var/log/shmlog/** directory. Files follow this naming convention: **<cfg-file-name>-<daemon>-<severity>.log**. The shmlog files are not human-readable, so to access the logs you must first run the following command to generate a file in the **/var/log/<file-name>/** directory with logs from all daemons:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<file-name>* logname all
```

If you then want to view logs from a specific daemon, you need to run the following command to generate a file under the **/var/log/<file-name>/** directory with complete logs:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<filename>  logname authd*
```

filtering Command to enable filtering. Filtering is subscriber centric and is useful for debugging and troubleshooting. It is disabled by default so you must use this option to enable it.

For example, if you want to quickly view the transmit packet logs for subscribers with interface-name pp0.100, you could use the following command to display only the relevant results:

```
user@rdevice> show shmlog entries logname jpppd* interface-name pp0.100 flag transmit-
packets
```

To debug sessions according to the interface name, use this command:

```
user@rdevice> show shmlog entries logname all interface-name pp0.100
```

To debug sessions that are logging in via VLAN 7 on physical-interface ge-0/0/0, use this command:

```
user@rdevice> show shmlog entries logname all vlan 7 physical-interface ge-0/0/0
```

log-name Name of the file containing the log output. Use this option to override all logs or a specified log, and to set the verbosity level (brief, detail, extensive, none, or terse). For example, to configure **bbe-autoconf-info** for detailed file logging, you would use the following command:

```
user@rdevice> [edit system services subscriber-management overrides shmlog]
user@rdevice> set log-name bbe-autoconf-info detail file-logging
```

log-type Severity level of the collected logs. Use this option to configure the severity level for captured logs (notice, info, or debug).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

short-cycle-protection (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 2382
- Hierarchy Level | 2382
- Description | 2383
- Options | 2383
- Required Privilege Level | 2383
- Release Information | 2383

Syntax

```
short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name]
[edit logical-systems name forwarding-options dhcp-relay ...],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay ...],
[edit routing-instances name forwarding-options dhcp-relay ...],
[edit logical-systems name system services dhcp-local-server ...],
[edit logical-systems name routing-instances name system services dhcp-local-server dhcp-local-
server...],
[edit routing-instances name system services dhcp-local-server ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
```

```
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name]
```

Description

Enable DHCP short-cycle protection to reduce resource usage associated with connection and authentication processing in highly scaled networks. You must configure both the minimum duration and the maximum duration for the lockout period.

The router detects short-lived client sessions and clients that repeatedly fail session negotiation, then locks them out from access by dropping subsequent DHCP discover or solicit messages from the client. The clients are tracked by the client identifier (client key), which can be a MAC address or some other unique value for DHCPv4 clients or the DUID for DHCPv6 clients. Locked-out clients are entered in the lockout database. If a locked-out client attempts another session before the grace time threshold is reached, it is locked out again. Each successive lockout period is increased exponentially up to the maximum lockout period. The grace time threshold is automatically set at whichever value is larger, 900 seconds or the configured maximum value.

Options

lockout-max-time *seconds*
: Maximum length of any lockout period; the upper bound of the lockout range.

- **Range:** 1 through 86400

lockout-min-time *seconds*
: Minimum length of any lockout period; the lower bound of the lockout period. The minimum value is the length of the first lockout period for a client. It cannot be greater than the maximum value. If you set it to the same value as the maximum, then the lockout period is fixed and does not increase for a client's subsequent lockouts.

- **Range:** 1 through 86400

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

DHCP Short Cycle Protection | 885

DHCP Short Cycle Protection | 885

smg-service (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 2384
- Hierarchy Level | 2384
- Description | 2385
- Required Privilege Level | 2385
- Release Information | 2385

Syntax

```
smg-service {
    failover other-routing-engine;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag <disable>;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Configure system services, including tracing operations and Routing Engine failover, for the main enhanced subscriber management session management process, `smg-service`.

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

trace—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Configuring the Subscriber Management Database Trace Log Filename

source-interface-set-at-login

IN THIS SECTION

- Syntax | **2386**
- Hierarchy Level | **2386**
- Description | **2386**
- Options | **2386**
- Required Privilege Level | **2386**
- Release Information | **2386**

Syntax

```
source-interface-set-at-login {
    svlan;
}
```

Hierarchy Level

```
[edit protocols ppp-service]
```

Description

Enable PPP to generate the name of an interface set for use during subscriber login. The AAA process then sends the name by means of the Juniper Networks VSA QoS-Set-Name (26-130) to the RADIUS server in the Access-Request message. You can use this statement to discriminate between business and residential subscribers in heterogeneous networks. The interface set name is created for business subscribers. The RADIUS server then returns the VSA in Access-Accept messages only for business subscribers.

Options

svlan Specifies that PPP constructs the name of an interface from the physical interface name and the outer (SVLAN) VLAN tag.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Automatic Creation of Business Subscriber Interface Sets

stacked-vlan-ranges (RADIUS Options)

IN THIS SECTION

- Syntax | 2387
- Hierarchy Level | 2387
- Description | 2387
- Options | 2388
- Required Privilege Level | 2388
- Release Information | 2388

Syntax

```
stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the stacked VLAN (S-VLAN) range of subscribers to which the named NAS-Port options definition applies.

NOTE: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

any	Entire S-VLAN range representing all S-VLAN IDs. The inner tag (S-VLAN ID) of the S-VLAN range must be configured as `any` to represent all inner VLAN ID tags.
low-outer-tag	Outer VLAN ID tag representing the lower limit of the S-VLAN range. • **Range:** 1 through 4094
high-outer-tag	Outer VLAN ID tag representing the upper limit of the S-VLAN range. • **Range:** 1 through 4094

NOTE: To specify a single outer VLAN ID tag, set *low-outer-tag* and *high-outer-tag* to the same value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **527**

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **526**

starts-with (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 2389
- Hierarchy Level | 2389
- Description | 2392
- Options | 2392
- Required Privilege Level | 2392
- Release Information | 2393

Syntax

```
equals {
    ascii name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    hexadecimal name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
```

```
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-60],
```

```
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
```

```
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Configure a partial match criteria used with the DHCP relay agent selective processing feature. DHCP relay agent compares the configured partial match string with the option-specific string received in DHCP client packets. If there is an partial left-to-right match, DHCP performs the action you define for the match criteria.

The option-specific string in the DHCP client packets can contain a superset of the specified ASCII or hexadecimal match string, provided that the leftmost characters of the option-specific string entirely match the characters in the configured match string.

You can configure an unlimited number of match strings. If you have multiple partial match configurations, the longest match rule applies. For example, DHCP relay agent matches the string “test123” before it matches the string “test”. Match strings do not support wildcard attributes.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

ascii-string	ASCII string of 1 through 255 alphanumeric characters.
hexadecimal-string	Hexadecimal string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Statement updated in Junos OS Release 17.4.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic | 727

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address | 747

static-subscribers (Dynamic Service Provisioning)

IN THIS SECTION

- Syntax | 2393
- Hierarchy Level | 2394
- Description | 2395
- Required Privilege Level | 2395
- Release Information | 2395

Syntax

```
static-subscribers {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            delimiter delimiter-character;
            domain-name domain-name;
            interface;
            logical-system-name;
            routing-instance-name;
```

```
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients  (merge | replace);
    }
    group group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                delimiter delimiter-character;
                domain-name domain-name;
                interface;
                logical-system-name;
                routing-instance-name;
                 user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
        }
        interface  interface-name <exclude> <upto upto-interface-name>;
    }
    service-profileservice-profile-name
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services],
[edit routing-instances routing-instances-name system services],
[edit system services]
```

Description

Configure and associate subscribers with statically configured interfaces for dynamic service provisioning.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview | 1487

statistics (Access Profile)

IN THIS SECTION

- Syntax | 2396
- Hierarchy Level | 2396
- Description | 2396
- Options | 2396
- Required Privilege Level | 2396
- Release Information | 2396

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options

`time`—Collect uptime statistics only.

`volume-time`—Collect both volume and uptime statistics.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`volume-time` option added in Junos OS Release 9.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

statistics (Service Accounting)

IN THIS SECTION

- Syntax | 2397
- Hierarchy Level | 2397
- Description | 2397
- Options | 2397
- Required Privilege Level | 2398
- Release Information | 2398

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name service accounting]
```

Description

Configure the router to collect time statistics, or both volume and time statistics, for the service accounting sessions being managed by AAA.

Options

`time`—Collect uptime statistics only.

`volume-time`—Collect both volume and uptime statistics.

NOTE: You must not configure the `volume-time` option if the `local` option is configured at the `[edit access profile profile-name service accounting-order]` hierarchy level, because the `local` configuration

automatically determines that both volume and time statistics are collected for the service accounting sessions. If you configure the `volume-time` option in this case, an error is generated when you commit the configuration.

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 14.2R1.

RELATED DOCUMENTATION

Configuring Service Accounting | 586

Configuring Per-Subscriber Session Accounting | 573

Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 583

strict (DHCP Local Server)

IN THIS SECTION

- Syntax | 2399
- Hierarchy Level | 2399
- Description | 2399
- Default | 2399
- Required Privilege Level | 2399
- Release Information | 2399

Syntax

```
strict;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Specify whether the server denies a client to bind when the client does not indicate that it accepts reconfigure messages. This feature is available only for DHCPv6.

Default

Accept solicit messages from clients that do not support reconfiguration and permit them to bind.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Preventing Binding of Clients That Do Not Support Reconfigure Messages | 912

strip-domain (Domain Map)

IN THIS SECTION

- Syntax | 2400
- Hierarchy Level | 2400
- Description | 2400
- Required Privilege Level | 2400
- Release Information | 2401

Syntax

```
strip-domain;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Remove the domain name from the username before continuing with any AAA services specified in a domain map.

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

Enabling Domain Name Stripping | 671

Configuring Domain and Realm Name Usage for Domain Maps | 667

strip-username (Domain Map)

IN THIS SECTION

- Syntax | 2401
- Hierarchy Level | 2401
- Description | 2402
- Options | 2402
- Required Privilege Level | 2402
- Release Information | 2402

Syntax

```
strip-username (left-to-right | right-to-left)
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Strip the user portion of the username for authentication to simplify off-chassis provisioning. Parsing direction can be left-to-right (default) or right-to-left.

Options

left-to-right Strip the user portion of the username from left to right.

right-to-left Strip the user portion of the username from right to left.

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 14.1.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning | 671

sub-domain

IN THIS SECTION

- Syntax | 2403
- Hierarchy Level | 2403
- Description | 2403
- Options | 2406
- Required Privilege Level | 2407
- Release Information | 2407

Syntax

```
sub-domain name {
 aaa-logical-system name {
      aaa-routing-instance (default | name);
 }
 target-logical-system name {
     target-routing-instance (default | name);
 }
 access-profile access-profile;
 address-pool address-pool;
 dynamic-profile dynamic-profile;
 override-chap-password override-chap-password;
 override-password override-password;
 qualifier {
  vlan-id-list [ vlan-id-list ... ];
 }
 strip-domain;
 strip-username (left-to-right | right-to-left);
 tunnel-profile tunnel-profile;
 using-user-password;
}
```

Hierarchy Level

```
[edit access domain map]
```

Description

You can configure subdomains under a domain map. Subdomain allows you to select different access-profile for users within the same domain different VLAN ID or VLAN ID range. Subdomain gives the flexibility to differentiate the users in a domain. It provides different services per the profile configuration.

The options available at the domain map level are also available at the subdomain. But the options defined for a subdomain are independent from other subdomains and the options at the domain map level.

Characteristics of a Subdomain

- Subdomain configuration within a domain takes the higher precedence than the domain map level configuration.
- The qualifier option is mandatory to define a subdomain.
- Qualifiers (VLAN ID) cannot have any overlap within a domain.
- You can configure maximum of 16 subdomains within a domain.

Here're some examples of different types of subdomain configurations and their behavior.

1. Access profile configuration in a domain with VLAN ID qualifier using subdomain.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub3 qualifier vlan-id-list 300-399
user@host# set sub-domain sub3 access-profile D
```

   This configuration creates different access profiles using the VLAN ID qualifier and subdomain.

2. Access profile and strip domain configuration through domain map.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 strip-domain
user@host# set sub-domain sub3 qualifier vlan-id-list 300-399
user@host# set sub-domain sub3 access-profile D
user@host# set sub-domain sub3 strip-domain
```

   The intention of this configuration is to use the domain map for an access profile selection, then strip the domain name. For each subdomain you can assign different access profile depending on the VLAN ID, but the `strip-domain` kept common for all cases. Though it appears a repeat configuration in

subdomain, since the subdomain is totally independent, it gives better flexibility when it comes to assign modifiers selectively.

3. Domain map and subdomain configuration for completely independent attribute selection.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 address-pool
user@host# set sub-domain sub3 qualifier vlan-id-list 900-399
user@host# set sub-domain sub3 access-profile D
user@host# set sub-domain sub3 dynamic-profile vlan-profile-9xx
```

For the same domain `abc.com`, each subdomain and unqualified domain map (top level) is independently defining its actions. VLAN ID is the qualifier for subdomain, which takes precedence and overrides the unqualified attributes with independent set available in the qualified subdomain.

4. Configure an empty subdomain.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list [ 100 200-299 400-450 ]
```

This configuration creates an empty subdomain with a set of VLAN ranges. This configuration is an example to exclude the users of the same domain depending on their VLAN ID.

Any user login qualifying for the subdomain match do not apply any options. All the other non-matching users in the domain get the options from un-qualified top level domain map.

5. Invalid subdomain configuration (overlapping VLAN ID ranges in subdomains).

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list [ 100 200-299 400-450 ]
user@host# set sub-domain sub1 access-profile B
```

```
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 250-300 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 strip-domain
```

This configuration gets rejected during commit. The qualifiers within the same domain map cannot have any overlap.

Example of an error message while trying such invalid subdomain configuration commit:

```
root@host# commit
2021-02-03 22:50:39.730422 IST: Running FIPS Self-tests
Veriexec is not enforced, FIPS mode not available
2021-02-03 22:50:39.768595 IST: FIPS Self-tests Skipped
[edit access domain map abc.com sub-domain sub2 qualifier]
  'vlan-id-list 250-300'
    Range 250-300 overlaps with range 200-299 in another sub-domain under same domain
error: configuration check-out failed
```

Options

sub-domain *name* Name of a subdomain.

aaa-logical-system Logical system used for applying AAA services.

aaa-routing-instance Routing instance used for applying AAA services.

- default—Default routing instance.
- name—Name of the routing instance you want to configure.

target-routing-instance Specify the routing instance of the subscriber context.

- default—Default routing instance.
- name—Name of the routing instance you want to configure.

access-profile *profile-name* Name of an access profile.

address-pool Specify the address pool used to assign addresses to subscribers associated with the domain map.

dynamic-profile	Specify the dynamic profile that is used for subscriber sessions associated with the subdomain.
override-chap-password	Use this CHAP password for authentication.
override-password	Use this password for authentication.
strip-domain	Enable domain name stripping from the username.
strip-username	Enable user name stripping from the username. • left-to-right—Strip to first domain delimiter on the left. • right-to-left—Strip to first domain delimiter on the right.
tunnel-profile	Specify the tunnel profile that provides definitions for tunnels associated with the subdomain.
using-user-password	Send overridden CHAP-Password using User-Password.

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 21.3R1.

RELATED DOCUMENTATION

Domain Mapping Overview | 655

subscriber (Access Profile)

IN THIS SECTION

- Syntax | 2408
- Hierarchy Level | 2408
- Description | 2408
- Options | 2409
- Required Privilege Level | 2410
- Release Information | 2410

Syntax

```
subscriber username {
    delegated-pool delegated-pool-name;
    framed-ip-address ipv4-address;
    framed-ipv6-pool ipv6-pool-name;
    framed-pool ipv4-pool-name;
    password password;
    target-logical-system logical-system-name<(target-routing-instance (default | routing-
instance-name)>;
    target-routing-instance (default | routing-instance-name);
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Enable local authentication for subscribers by configuring a password to match the subscriber. Local authentication can take the form of either user password authentication or Challenge Handshake Authentication Protocol' (CHAP) authentication. For user password authentication, the configured password is used to verify the subscriber's login password. For CHAP authentication, the configured

password acts as the challenge secret to verify the subscriber's challenge password and challenge response credential.

NOTE: Local authentication and authorization also requires the `password` option to be configured as an `authentication-order` method for the access profile.

You can also optionally configure several attributes, such as an address, address pool, logical system, or routing instance, to be authorized locally for the subscriber when authentication is successful.

Local authentication supports all subscriber types that are currently supported by subscriber management and services on MX Series routers.

Local authentication is useful when you do not want to use external authentication servers. The associated local authorization similarly is useful when you do not want to use external authorization servers. Another use case might be when you are migrating a network from E Series routers running JunosE software to MX Series routers running Junos OS. You may also want to configure local authentication and authorization as a backup for RADIUS authentication.

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 supports 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 subscriber count against the system limit of 100 subscribers for local authentication.

Options

delegated-pool *delegated-pool-name*
(Optional) Specify the name of an address pool used to locally allocate a delegated IPv6 prefix for the subscriber. Corresponds to RADIUS standard attribute Delegated-IPv6-Prefix (123).

framed-ip-address *ipv4-address*
(Optional) Specify the IP address to be configured for the subscriber. Corresponds to RADIUS standard attribute Framed-IP-Address (8).

framed-ipv6-pool *ipv6-pool-name*
(Optional) Specify the name of an address pool used to assign a router advertisement IPv6 prefix or a DHCPv6 IA_NA/128 address for the subscriber. Corresponds to RADIUS standard attribute Framed-IPv6-Pool (100).

framed-pool *ipv4-pool-name*	(Optional) Specify the name of an address pool used to assign an IPv4 address for the subscriber. Corresponds to RADIUS standard attribute Framed-Pool (88).
password *password*	Specify the password used to authenticate the subscriber locally. Corresponds to RADIUS standard attributes User-Password (2) or CHAP-Password (3).
target-logical-system *logical-system-name*	(Optional) Specify the name of the logical system assigned to the subscriber.
target-routing-instance (default | *routing-instance-name*)	(Optional) Specify the name of the routing instance assigned to the subscriber; either the default routing instance or a nondefault routing instance.

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 18.2R1.

RELATED DOCUMENTATION

Example: Configuring CHAP Authentication with RADIUS

Specifying the Authentication and Accounting Methods for Subscriber Access | 551

Configuring Local Authentication and Authorization for Subscribers | 552

subscriber-packet-idle-timeout

IN THIS SECTION

- Syntax | 2411
- Hierarchy Level | 2411
- Description | 2411

Options | 2411
Required Privilege Level | 2411

Syntax

```
subscriber-packet-idle-timeout subscriber-packet-idle-timeout;
```

Hierarchy Level

```
[edit system services packet-triggered-subscribers]
```

Description

The subscriber packet idle timeout for packet triggered subscribers.

Options

subscriber-packet-idle-timeout—Maximum idle time.

- **Range:** 15 through 1440 minutes.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Packet-Triggered Subscribers Services Overview

subscriber-management (Subscriber Management)

IN THIS SECTION

- Syntax | 2412
- Hierarchy Level | 2413
- Description | 2413
- Required Privilege Level | 2414
- Release Information | 2414

Syntax

```
subscriber-management {
    enable;
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    }
    overrides {
        event {
            catastrophic-failure {
                reboot (master | standby);
            }
        }
        interfaces {
            family (inet | inet6) {
                layer2-liveness-detection;
                ipoe-dynamic-arp-enable;
                receive-gratuitous-arp;
            }
        }
        no-unsolicited-ra;
        ra-initial-interval-max seconds;
        ra-initial-interval-min seconds;
        shmlog {
            disable;
            file filename <files maximum-no-files> <size maximum-file-size—>;
            filtering enable;
```

```
            log-name {
                all;
                logname {
                    <brief | detail | extensive | none | terse>;
                    <file-logging |no-file-logging>;
                }
                }
            log-type (debug | info | notice);
        |
    }
    redundancy {
        interface name {
            local-inet-address v4-address;
            local-inet6-address v6-address;
            shared-key string;
            virtual-inet-address virtual-v4-address;
            virtual-inet6-address virtual-v6-address;
        }
        no-advertise-routes-on-backup;
        protocol {
            pseudo-wire;
            vrrp;
        }
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Maintaining DHCP Subscribers During Interface Delete Events | 865
- *Tracing Subscriber Management Database Events for Troubleshooting*
- *Junos OS Enhanced Subscriber Management*
- *Configuring Junos OS Enhanced Subscriber Management*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets
- *Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover*
- M:N Subscriber Redundancy on BNG | 1173

subscriber-profile

IN THIS SECTION

- Syntax | 2415
- Hierarchy Level | 2415
- Description | 2415
- Options | 2415
- Required Privilege Level | 2415
- Release Information | 2415

Syntax

```
subscriber-profile profile-name {
    enable service-name {
        concurrent-data-sessions max-session-number;
    }
    disable service-name;
    max-data-sessions-per-subscriber {
        limit max-sub-sessions;
        exceed-action {
            drop;
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit services service-set services-set-name]
```

Description

Specify the subscriber profile name. A subscriber profile specifies which services should be enabled and which services should be disabled for traffic belonging to a subscriber bound to a particular subscriber profile. A subscriber is bound to a minimum of one subscriber profile at any given time.

Options

profile-name—Name of the 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 11.4.

subscription-id-data-include (PCRF Partition)

IN THIS SECTION

- Syntax | 2416
- Hierarchy Level | 2416
- Description | 2416
- Options | 2417
- Required Privilege Level | 2418
- Release Information | 2418

Syntax

```
subscription-id-data-include {
    base-interface-name;
    delimiter delimiter-character;
    domain-name name;
    interface-name;
    mac-address;
    nas-port-id;
    origin-host;
    origin-realm;
    user-name;
    user-prefix prefix;
    vlan-tags;
}
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the subscriber data to use in a Policy and Charging Rules Function (PCRF).

Options

base-interface-name Use the specified physical or underlying interface name. If both the underlying interface and the client application perform authentication, authorization, and provisioning, the identification attributes in the server requests enable the AAA or PCRF server to make an association between the two entities.

delimiter *delimiter-character* Use the specified character used as the delimiter between the concatenated components of the subscription-id-data-include. You cannot use the semicolon (;) as a delimiter.

- **Default:** @

domain-name *name* Use the specified domain name that is concatenated with the subscription-id-data-include during the subscriber identification process.

interface-name Use the specified interface name that is concatenated with the subscription-id-data-include during the subscriber identification process; for example, demux0.

mac-address Use the specified client hardware address (chaddr) from the incoming packet that is concatenated with the subscription-id-data-include during the subscriber identification process.

nas-port-id Use the specified NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface that subscriber management uses to identify subscribers that is concatenated with the subscription-id-data-include during the subscriber identification process. By default, the NAS-Port-ID includes the interface-description value that describes the physical interface.

origin-host Use the specified name of the host that originates the Diameter message that is concatenated with the subscription-id-data-include during the subscriber identification process. Supplied as the value of Origin-Host AVP for all messages sent by the Diameter master instance.

origin-realm Use the specified realm of the host that originates the Diameter message that is concatenated with the subscription-id-data-include during the subscriber identification process. Supplied as the value of Origin-Realm AVP for all messages sent by the Diameter master instance.

user-name Use the specified subscriber username that is concatenated with the subscription-id-data-include during the subscriber identification process.

user-prefix *prefix* Use the specified user prefix that is concatenated with the subscription-id-data-include during the subscriber identification process.

vlan-tags Include the subscriber session VLAN tags in the subscription ID. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

Use this option instead of the `interface-name` option when you 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.

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 16.2.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

- Configuring the PCRF Partition | **1455**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gx Interactions Between the Router and the PCRF | **1418**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**
- Understanding Upstream and Downstream Messages for the PCRF | **1444**

subscription-id-type (PCRF Partition)

IN THIS SECTION

- Syntax | **2419**
- Hierarchy Level | **2419**

- Description | 2419
- Options | 2419
- Required Privilege Level | 2420
- Release Information | 2420

Syntax

```
subscription-id-type number;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Identify the subscriber. You can identify the subscriber using several IANA predefined values, or define your own. The Policy and Charging Rule Function (PCRF) uses the subscription-id-type during a subscriber login session using for CCR-GX-I and CCA-GX-I messages.

BEST PRACTICE: For wireline service, we recommend that you define your own customer-specific value.

If you do not define or include any data in the `subscription-id-data-include` statement, then the value defaults to reserved. For some applications, the PCRF may be configured to identify subscribers using another method, and instead use the subscription-id-type for protocol compliance. If the dynamic-subscription-id is provided by the AAA server, then the PCRF uses the dynamic-subscription-id as the identifier regardless of what you defined for the local configuration.

Options

number	Identifier for the type of subscriber for a PCRF partition.

- **Default:** 4 (END_USER_PRIVATE)
- **Range:** 1 through 2147483647
- **Values:** 1 (END_USER_IMSI)
  2 (END_USER_SIP_URI)
  3 (END_USER_NAI)
  4 (END_USER_PRIVATE)

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | **1455**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | **1418**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

Understanding Upstream and Downstream Messages for the PCRF | **1444**

target-logical-system (Domain Map)

IN THIS SECTION

- Syntax | **2421**
- Hierarchy Level | **2421**
- Description | **2421**
- Default | **2421**

- Options | 2421
- Required Privilege Level | 2422
- Release Information | 2422

Syntax

```
target-logical-system logical-system-name {
    target-routing-instance routing-instance-name;
}
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Configure a non-default logical system and optionally a non-default routing instance for the subscriber's interface in a domain map.

You use the `target-routing-instance` statement at the `[edit access domain map domain-map-name]` hierarchy level to configure a non-default routing instance for the default logical system.

NOTE: Subscriber management is supported in the default logical system only.

Default

Default logical system for the subscriber..

Options

logical-system-name—Name of the logical system.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying a Target Logical System/Routing Instance in a Domain Map | 665

target-routing-instance (Domain Map)

IN THIS SECTION

- Syntax | 2422
- Hierarchy Level | 2423
- Description | 2423
- Default | 2423
- Options | 2423
- Required Privilege Level | 2423
- Release Information | 2423

Syntax

```
target-routing-instance (routing-instance-name | default);
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit access domain map domain-map-name target-logical-system logical-system-name]
```

Description

Configure the routing instance of the subscriber context.

NOTE: Subscriber management is supported in the default logical system only. The `target-logical-system` statement, which appears in the CLI, is not supported in current Junos OS releases.

Default

For dynamic LNS sessions, the routing instance of the peer (LAC facing) interface. For all other sessions, the default logical system/routing instance context.

Options

routing-instance-name—Name of the routing instance.

`default`—The default routing instance.

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.

`default` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Specifying a Target Logical System/Routing Instance in a Domain Map | 665

Domain Mapping Overview | 655

terminate-code

IN THIS SECTION

- Syntax | 2424
- Hierarchy Level | 2424
- Description | 2424
- Options | 2425
- Required Privilege Level | 2426
- Release Information | 2426

Syntax

```
terminate-code (aaa (deny | service-shutdown | shutdown) | dhcp | l2tp | ppp | vlan) term-reason
radius value;
```

Hierarchy Level

```
[edit access]
```

Description

Customize the mapping between a termination cause (the internal termination identifier) and a numerical code value for the cause that is reported in the RADIUS Acct-Terminate-Cause attribute (49).

When a RADIUS Acct-Stop message is issued as a result of the termination of a subscriber 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 for AAA, DHCP, L2TP, PPP, and VLAN subscriber and service session failures. By default, these internal cause codes 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 the events.

Because there are many different Junos OS internal identifiers for termination causes and only 18 supported, RFC-defined standard code values, by default a given code value can map to multiple identifiers. Instead of using the default code values, you can use the `terminate-code` statement to 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.

Options

aaa — Map internal identifiers for AAA-specific termination causes to a numerical value.

deny — Limit selection of termination causes to those associated with denial of subscriber access.

dhcp — Map internal identifiers for DHCP-specific termination causes to a numerical value.

l2tp — Map internal identifiers for L2TP-specific termination causes to a numerical value.

ppp — Map internal identifiers for PPP-specific termination causes to a numerical value.

radius *value* — Number that represents the termination cause in the RADIUS Acct-Terminate-Cause attribute (49).

- **Range:** 1 through 4,294,967,295

service-shutdown — Limit selection of termination causes to those associated with established service sessions independent of the parent subscriber session.

shutdown — Limit selection of termination causes to those associated with established subscriber sessions.

vlan — Map internal identifiers for VLAN-specific termination causes to a numerical value.

term-reason — Internal identifier for the termination causes defined for the specified protocol type.

For protocol-specific termination causes, see the following topics:

- "AAA Termination Causes and Code Values" on page 608

- "DHCP Termination Causes and Code Values" on page 610
- "L2TP Termination Causes and Code Values" on page 611
- "PPP Termination Causes and Code Values" on page 638
- "VLAN Termination Causes and Code Values" on page 651

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`vlan` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Session Termination Causes and RADIUS Termination Cause Codes | 603

Mapping Session Termination Causes to Custom Termination Cause Codes | 606

timeout (DHCP Local Server)

IN THIS SECTION

- Syntax | 2427
- Hierarchy Level | 2427
- Description | 2427
- Options | 2428
- Required Privilege Level | 2428
- Release Information | 2428

Syntax

```
timeout timeout-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure the initial value in seconds between attempts to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.

Options

timeout-value—Initial retry timeout value.

- **Range:** 1 through 10 seconds
- **Default:** 2 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Configuring Dynamic Reconfiguration Attempts for DHCP Clients | 875

timeout-grace (Access)

IN THIS SECTION

- Syntax | 2429
- Hierarchy Level | 2429
- Description | 2429
- Options | 2429
- Required Privilege Level | 2429
- Release Information | 2429

Syntax

```
timeout grace seconds;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure a grace period during which a RADIUS server that times out after failing to respond to an authentication request is not considered to be down (if other servers are available) or unreachable (if it is the only configured server). The server is marked as down or unreachable only when it does not successfully respond to an authentication request during the grace period and times out again after the period expires.

You can use long grace periods to give unresponsive servers more opportunities to respond before abandoning them. You can use short grace periods to cause the router to declare unresponsive servers down and direct requests to available servers sooner.

Options

seconds—Duration of the server timeout grace period.

- **Range:** 0 through 30
- **Default:** 10

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 17.2.

RELATED DOCUMENTATION

Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable | 482

RADIUS Servers and Parameters for Subscriber Access | 476

token (DHCP Local Server)

IN THIS SECTION

- Syntax | 2430
- Hierarchy Level | 2430
- Description | 2431
- Options | 2431
- Required Privilege Level | 2431
- Release Information | 2431

Syntax

```
token token-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
```

```
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure a plain-text token for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. The token enables rudimentary entity authentication to protect against inadvertently instantiated DHCP servers. A null token (empty string) indicates that the configuration token functionality is not enabled. A group configuration takes precedence over a DHCP local server configuration. For more information about tokens, see RFC 3118, *Authentication for DHCP Messages*, section 4.

Options

token-value—Plain-text alphanumeric string.

- **Default:** null (empty string)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

Configuring a Token for DHCP Local Server Authentication | 877

trace (DHCP Local Server)

IN THIS SECTION

- Syntax | 2432
- Hierarchy Level | 2432
- Description | 2433
- Required Privilege Level | 2433
- Release Information | 2433

Syntax

```
trace;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
```

```
group-name interface interface-name],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
interface interface-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name interface interface-name]
```

Description

Enable trace operations for a group of interfaces or for a specific interface within a group.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Tracing Extended DHCP Operations | 902

Tracing Extended DHCP Operations | 902

trace (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2434
- Hierarchy Level | 2434
- Description | 2434
- Required Privilege Level | 2434
- Release Information | 2434

Syntax

```
trace;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name  interface-name],
[edit  logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
interface interface-name],
[edit  logical-systems logical-system-name forwarding-options dhcp-relay group group-name
interface interface-name],
[edit  logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit  logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name interface interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name interface interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name]
```

Description

Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Configuring DHCP Relay Agent
- DHCP Monitoring and Management | 895

traceoptions (ANCP)

IN THIS SECTION

- Syntax | 2435
- Hierarchy Level | 2435
- Description | 2435
- Options | 2436
- Required Privilege Level | 2437
- Release Information | 2437

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag <disable>;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Define tracing operations for ANCP agent processes.

Options

`file` *`filename`*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. Include the `disable` option after a flag to disable tracing for that flag. You can include the following flags:

- `all`—Trace all operations.
- `config`—Trace configuration events.
- `cos`—Trace class-of-service events.
- `general`—Trace general flow.
- `packet`—Trace ANCP packet transmit and receive operations.
- `process`—Trace process internals.
- `protocol`—Trace protocol events.
- `restart`—Trace process restart flow
- `routing-socket`—Trace routing socket events.
- `session`—Trace connection events and flow.
- `startup`—Trace ANCP startup events and flow.
- `subscriber`—Trace subscriber events.
- `timer`—Trace timer processing.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.

- `notice`—Match notice messages about conditions reqcliring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Tracing ANCP Events for Troubleshooting | 1333

traceoptions (DHCP)

IN THIS SECTION

- Syntax | 2438
- Hierarchy Level | 2438
- Description | 2438
- Options | 2439
- Required Privilege Level | 2441
- Release Information | 2441

Syntax

```
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;
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
[edit security dynamic-address]
```

Description

Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.

This statement replaces the deprecated `traceoptions` statements at the `[edit forwarding-options dhcp-relay]` and `[edit system services dhcp-local-server]` hierarchy levels.

NOTE: Traceoptions does not differentiate between a logical system and tenant system, and can be configured under the root logical system.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Trace all events.
- `auth`—Trace authentication events.
- `database`—Trace database events.
- `fwd`—Trace firewall process events.
- `general`—Trace miscellaneous events.
- `ha`—Trace high availability-related events.
- `era`—Enables logging.
- `interface`—Trace interface operations.
- `io`—Trace I/O operations.
- `liveness-detection`—Trace liveness detection operations.
- `packet`—Trace packet and option decoding operations.
- `performance`—Trace performance measurement operations.
- `profile`—Trace profile operations.
- `rpd`—Trace routing protocol process events.

- `rtsock`—Trace routing socket operations.
- `security-persistence`—Trace security persistence events.
- `session-db`—Trace session database events.
- `state`—Trace changes in state.
- `statistics`—Trace baseline statistics.
- `ui`—Trace user interface operations.

`level`—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access, allowing only the user `root` and users who have the Junos OS `maintenance` permission to access the trace files.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (*`maximum-file-size`*`k`), megabytes (*`maximum-file-size`*`m`), or gigabytes (*`maximum-file-size`*`g`). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10,240 through 1,073,741,824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

traceoptions (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 2441
- Hierarchy Level | 2442
- Description | 2442
- Options | 2442
- Required Privilege Level | 2443
- Release Information | 2443

Syntax

```
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;
}
```

Hierarchy Level

```
[edit system processes diameter-service]
```

Description

Define tracing options for Diameter processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations
- `application`—Trace Diameter application interface events
- `configuration`—Trace configuration events
- `daemon`—Trace Diameter daemon level events
- `diameter-instance`—Trace Diameter instance events
- `dne`—Trace Diameter network element events
- `framework`—Trace Diameter framework events
- `memory-management`—Trace memory management events
- `messages`—Trace Diameter messages

- `node`—Trace Diameter node events
- `peer`—Trace Diameter peer events

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Tracing Diameter Base Protocol Processes for Subscriber Access

traceoptions (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | 2444
- Hierarchy Level | 2444
- Description | 2444
- Options | 2444
- Required Privilege Level | 2445
- Release Information | 2445

Syntax

```
traceoptions file file-name flag authentication | flag configuration | flag dictionary | flag
fsm | flag statistics | flag kernel | flag dynamic | flag database | flag op-script | flag
general | flag all
```

Hierarchy Level

```
[edit system processes extensible-subscriber-services]
```

Description

Configure the trace options for Extensible Subscriber Services Manager. essmd supports Junos OS daemon trace options.

Options

traceoptions Configure the trace options for Extensible Subscriber Services Manager.

file-name	Name of the trace file.
flag all	Configure trace settings for all operations.
flag authentication	Configure trace settings for authentication operations.
flag configuration	Configure trace settings for configuration operations.
flag database	Configure trace settings for database operations.
flag dictionary	Configure trace settings for dictionary operations.
flag dynamic	Configure trace settings for dynamic profile operations.
flag fsm	Configure trace settings for finite state machine operations.
flag general	Configure trace settings for general operations.
flag kernel	Configure trace settings for kernel state-change operations.
flag op-script	Configure trace settings for op script operations.
flag statistics	Configure trace settings for statistics-collection operations.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

traceoptions (Enhanced Subscriber Management)

IN THIS SECTION

Syntax | 2446

- Hierarchy Level | 2446
- Description | 2446
- Options | 2446
- Required Privilege Level | 2448
- Release Information | 2448

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag <disable>;
    level level;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes smg-service]
```

Description

Define tracing operations for the main enhanced subscriber management session management process.

Options

`file` *`filename`*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All trace logs are stored in the directory `/var/log/bbesmgd`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 5 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. Include the `disable` option after a flag to disable tracing for that flag. You can include the following flags:

- `all`—Trace all operations.
- `auth`—Trace general authentication events.
- `autoconf`—Trace autoconfiguration events.
- `config`—Trace configuration events.
- `demux`—Trace demultiplexing (demux) events.
- `dhcp`—Trace DHCP Server events.
- `dprof`—Trace dynamic profile events.
- `interface`—Trace interface events.
- `io`—Trace packet socket events.
- `l2tp`—Trace L2TP events.
- `main`—Trace main operations.
- `net`—Trace network processing events.
- `ppp`—Trace PPP events.
- `pppoe`—Trace PPPoE events.
- `rpd`—Trace routing protocol process events.
- `rtsock`—Trace routing socket events.
- `service`—Trace service events.
- `session`—Trace connection events and flow.
- `stats`—Trace statistics events.
- `ucac`—Trace universal call admission control events.
- `ui`—Trace user interface events.
- `vbf`—Trace variable-based forwarding events.

`level` *`level`*—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions reqcliring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

> *Configuring Junos OS Enhanced Subscriber Management*
>
> *Configuring the Subscriber Management Database Trace Log Filename*

traceoptions (General Authentication Service)

IN THIS SECTION

- Syntax | 2449
- Hierarchy Level | 2449
- Description | 2449
- Options | 2450
- Required Privilege Level | 2451
- Release Information | 2451

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        user user@domain;
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes general-authentication-service]
```

Description

Configure tracing options for the general authentication service.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`filter`—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifier simplifies troubleshooting in a scaled environment.

- `user` *user@domain*—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `address-assignment`—Trace address-assignment pool events
- `all`—Trace all tracing operations
- `configuration`—Trace configuration events
- `framework`—Trace authentication framework events
- `gx-plus`—Trace Gx-Plus events
- `jsrc`—Trace JSRC events
- `ldap`—Trace LDAP authentication events
- `local-authentication`—Trace local authentication events
- `radius`—Trace RADIUS authentication events
- `user-access`—Trace user access events, such as login, logout, and authenticate.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB),

megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*`k` to specify KB, *size*`m` to specify MB, or *size*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Tracing General Authentication Service Processes

traceoptions (Static Subscribers)

IN THIS SECTION

- Syntax | 2452
- Hierarchy Level | 2452
- Description | 2452
- Options | 2452
- Required Privilege Level | 2454
- Release Information | 2454

Syntax

```
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;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system processes static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
processes static-subscribers],
[edit routing-instances routing-instances-name system processes static-subscribers],
[edit system processes static-subscribers]
```

Description

Define tracing operations for static subscriber processes.

Options

`file` *filename*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `authentication`—Trace authentication events.

- `configuration`—Trace configuration events.
- `database`—Trace database events.
- `general`—Trace general events.
- `gres`—Trace GRES events.
- `profile`—Trace dynamic profile events.
- `rtsock`—Trace routing socket events.
- `statistics`—Trace statistics events.
- `subscriber`—Trace subscriber events.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—(Optional) Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Tracing Static Subscriber Events for Troubleshooting | 1514

transport (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 2454
- Hierarchy Level | 2455
- Description | 2455
- Options | 2455
- Required Privilege Level | 2455
- Release Information | 2455

Syntax

```
transport transport-name {
    address;
    logical-system logical-system-name <routing-instance routing-instance-name >;
    routing-instance routing-instance-name
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Configure the Diameter instance and the local IP address for the Diameter local transport connection.

Options

transport-name—Name of the transport.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Diameter | 1373

Configuring the Diameter Transport | 1376

transport (Diameter Peer)

IN THIS SECTION

Syntax | 2456

- Hierarchy Level | 2456
- Description | 2456
- Default | 2456
- Options | 2456
- Required Privilege Level | 2456
- Release Information | 2457

Syntax

```
transport transport-name;
```

Hierarchy Level

```
[edit diameter peer peer-name connect-actively]
```

Description

Specify the transport layer connection to be used for establishing active connections to the peer.

Default

The transport is defined in the default logical system and master routing instance.

Options

transport-name—Name of the transport.

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 11.2.

RELATED DOCUMENTATION

- Configuring Diameter | 1373
- Configuring Diameter Peers | 1374

traps

IN THIS SECTION

- Syntax | 2457
- Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series) | 2457
- Hierarchy Level (QFX Series, EX4600) | 2458
- Description | 2458
- Required Privilege Level | 2458
- Release Information | 2458

Syntax

```
(traps | no-traps);
```

Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series)

```
[edit dynamic-profiles profile-name interfaces interface-name],
[edit interfaces interface-name],
[edit interfaces interface-name  unit logical-unit-number],
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name  unit logical-unit-number]
```

Hierarchy Level (QFX Series, EX4600)

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range interface-range-name]
```

Description

Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

(Enhanced subscriber management for MX Series routers) To enable SNMP notifications, you must first configure the `interface-mib` statement at the `[edit dynamic-profiles profile-name interfaces interface-name]` hierarchy level. If `interface-mib` is not configured, the `traps` statement has no effect.

BEST PRACTICE: To achieve maximum performance when enhanced subscriber management is enabled, we recommend that you *not* enable SNMP notifications on all dynamic subscriber interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Support at the `[edit dynamic-profiles profile-name interfaces interface-name]` hierarchy level introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

RELATED DOCUMENTATION

Enabling or Disabling SNMP Notifications on Physical Interfaces

Enabling or Disabling SNMP Notifications on Logical Interfaces

trigger (DHCP Local Server)

IN THIS SECTION

- Syntax | 2459
- Hierarchy Level | 2459
- Description | 2460
- Required Privilege Level | 2460
- Release Information | 2460

Syntax

```
trigger {
    radius-disconnect;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
```

```
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure behavior in response to a trigger for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | **874**

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect | **876**

radius-disconnect (DHCP Local Server) | **2244**

trio-flow-offload

IN THIS SECTION

- Syntax | 2461
- Hierarchy Level | 2461
- Description | 2461
- Options | 2462
- Required Privilege Level | 2462
- Release Information | 2462

Syntax

```
trio-flow-offload minimum-bytes minimum-bytes;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Enable any plug-in or daemon on a PIC to generate a request to off-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs).

NOTE: This feature is not supported for Broadband Edge subscribers (given that service PIC off load is not available with aggregate Ethernet (AE).

Options

minimum-bytes—Minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Flow Offloading on MX Series Routers

trust-option-82

IN THIS SECTION

- Syntax | 2462
- Hierarchy Level | 2463
- Description | 2463
- Required Privilege Level | 2463
- Release Information | 2463

Syntax

```
trust-option-82;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Enable processing of DHCP client packets that have a gateway IP address (giaddr) of 0 (zero) and contain option 82 information. By default, the DHCP relay agent treats such packets as if they originated at an untrusted source, and drops them without further processing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Enable Processing of Untrusted Packets So Option 82 Information Can Be Used | 761

Overriding the Default DHCP Relay Configuration Settings | 709

tunnel-profile (Domain Map)

IN THIS SECTION

- Syntax | 2464
- Hierarchy Level | 2464
- Description | 2464
- Options | 2464
- Required Privilege Level | 2464
- Release Information | 2465

Syntax

```
tunnel-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Tunnel profile that provides definitions for tunnels associated with the domain map.

Options

profile-name—Name of tunnel 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 Tunnel Profile in a Domain Map | 666

Configuring a Tunnel Profile for Subscriber Access

underlying-interface (ANCP)

IN THIS SECTION

- Syntax | 2465
- Hierarchy Level | 2465
- Description | 2466
- Options | 2466
- Required Privilege Level | 2466
- Release Information | 2466

Syntax

```
underlying-interface underlying-interface-name;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-set interface-set-name]
```

Description

Configure the underlying interface on which the VLAN demux interface is running. The VLAN demux interface is the underlying interface for the PPPoE sessions controlled by ANCP.

Options

underlying-interface-name	Name of the underlying interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring the ANCP Agent | 1253

Associating an Access Node with Subscribers for ANCP Agent Operations | 1256

unique-nas-port (Access)

IN THIS SECTION

- Syntax | 2467
- Hierarchy Level | 2467
- Description | 2467
- Options | 2467
- Required Privilege Level | 2467
- Release Information | 2468

Syntax

```
unique-nas-port {
    chassis-id chassis-id;
    chassis-id-width chassis-id-width;
}
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the router to provide a unique value for the NAS-Port attribute (RADIUS attribute 5) of each subscriber. You can configure a NAS-Port value that is unique within the router only, or unique across the different MX routers in the network.

Options

chassis-id Value for the `chassis-id` part of NAS-Port attribute. To ensure that NAS-Port values are unique across all MX series routers in the network, you must specify a unique `chassis-id` for each MX router.

- **Range:** 0-127 bits

chassis-id-width Number of bits used to uniquely identify the chassis across the MX series routers in the network. All routers must use the same `chassis-id-width`. If you do not configure a `chassis-id-width`, the resulting NAS-Port attribute is unique within the router only.

- **Range:** 1-7 bits

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 15.1.

RELATED DOCUMENTATION

Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers | 523

RADIUS Servers and Parameters for Subscriber Access | 476

unit

IN THIS SECTION

- Syntax | 2468
- Hierarchy Level | 2478
- Description | 2478
- Options | 2478
- Required Privilege Level | 2478
- Release Information | 2478

Syntax

```
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
```

```
        upstream-rate rate;
    }
    allow-any-vci;
    atm-scheduler-map (map-name | default);
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;

            }
            dynamic-profile profile-name;
        }
    }
    backup-options {
        interface interface-name;
    }
    bandwidth rate;
    cell-bundle-size cells;
    clear-dont-fragment-bit;
    compression {
        rtp {
            maximum-contexts number <force>;
            f-max-period number;
            queues [queue-numbers];
            port {
                minimum port-number;
                maximum port-number;
            }
        }
    }
    compression-device interface-name;
    copy-tos-to-outer-ip-header;
    demux {
        inet {
            address-source address;
            auto-configure {
                address-ranges {
                    authentication {
```

```
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }

        }
    }
    inet6 {
        address-source address;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
```

```
                    }
                }
            }

        }
    }
}
demux-destination family;
demux-source family;
demux-options {
    underlying-interface interface-name;
}
description text;
etree-ac-role (leaf | root);
interface {
    l2tp-interface-id name;
    (dedicated | shared);
}
dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [dial-string-numbers];
    idle-timeout seconds;
    incoming-map {
        caller caller-id | accept-all;
        initial-route-check seconds;
        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        watch-list {
            [routes];
        }
    }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
```

```
        bearer-bandwidth-limit kilobits-per-second;
    }
    encapsulation type;
    epd-threshold cells plp1 cells;
    family family-name {
        ... the family subhierarchy appears after the main [edit interfaces interface-name unit
logical-unit-number] hierarchy ...
    }
    fragment-threshold bytes;
    host-prefix-only;
    inner-vlan-id-range start start-id end end-id;
    input-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap |
        swap-push | swap-swap);
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        tag-protocol-id tpid;
        vlan-id number;
    }
    interleave-fragments;
    inverse-arp;
    layer2-policer {
        input-policer policer-name;
        input-three-color policer-name;
        output-policer policer-name;
        output-three-color policer-name;
    }
    link-layer-overhead percent;
    minimum-links number;
    mrru bytes;
    multicast-dlci dlci-identifier;
    multicast-vci vpi-identifier.vci-identifier;
    multilink-max-classes number;
    multipoint;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
    oam-period (disable | seconds);
    output-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap |
        swap-push | swap-swap);
        inner-tag-protocol-id tpid;
```

```
        inner-vlan-id number;
        tag-protocol-id tpid;
    }
    passive-monitor-mode;
    peer-unit unit-number;
    plp-to-clp;
    point-to-point;
    ppp-options {
        mru size;
        mtu (size | use-lower-layer);
        chap {
            access-profile name;
            default-chap-secret name;
            local-name name;
            passive;
        }
        compression {
            acfc;
            pfc;
        }
        dynamic-profile profile-name;
        ipcp-suggest-dns-option;
        lcp-restart-timer milliseconds;
        loopback-clear-timer seconds;
        ncp-restart-timer milliseconds;
        pap {
            access-profile name;
            default-pap-password password;
            local-name name;
            local-password password;
            passive;
        }
    }
    pppoe-options {
        access-concentrator name;
        auto-reconnect seconds;
        (client | server);
        service-name name;
        underlying-interface interface-name;
    }
    pppoe-underlying-options {
        access-concentrator name;
        direct-connect;
```

```
        dynamic-profile profile-name;
        max-sessions number;
    }
    proxy-arp;
    service-domain (inside | outside);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
burst length);
        queue-length number;
    }
    short-sequence;
    targeted-distribution;
    transmit-weight number;
    (traps | no-traps);
    trunk-bandwidth rate;
    trunk-id number;
    tunnel {
        backup-destination address;
        destination address;
        key number;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
    vci vpi-identifier.vci-identifier;
    vci-range start start-vci end end-vci;
    vpi vpi-identifier;
    vlan-id number;
    vlan-id-range number-number;
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    family family {
        accounting {
            destination-class-usage;
            source-class-usage {
                (input | output | input output);
            }
        }
        access-concentrator name;
        address address {
            ... the address subhierarchy appears after the main [edit interfaces interface-name
unit logical-unit-number family family-name] hierarchy ...
```

```
        }
        bundle interface-name;
        core-facing;
        demux-destination {
            destination-prefix;
        }
        demux-source {
            source-prefix;
        }
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
            group filter-group-number;
            input filter-name;
            input-list [filter-names];
            output filter-name;
            output-list [filter-names];
        }
        interface-mode (access | trunk);
        ipsec-sa sa-name;
        keep-address-and-control;
        mac-validate (loose | strict);
        max-sessions number;
        mtu bytes;
        multicast-only;
        no-redirects;
        policer {
            arp policer-template-name;
            input policer-template-name;
            output policer-template-name;
        }
        primary;
        protocols [inet iso mpls];
        proxy inet-address address;
        receive-options-packets;
        receive-ttl-exceeded;
        remote (inet-address address | mac-address address);
        rpf-check {
            fail-filter filter-name
            mode loose;
        }
        sampling {
```

```
            input;
            output;
        }
        service {
            input {
                post-service-filter filter-name;
                service-set service-set-name <service-filter filter-name>;
            }
            output {
                service-set service-set-name <service-filter filter-name>;
            }
        }
        service-name-table table-name
        targeted-options {
            backup backup;
            group group;
            primary primary;
            weight ($junos-interface-target-weight | weight-value);
        }
        (translate-discard-eligible | no-translate-discard-eligible);
        (translate-fecn-and-becn | no-translate-fecn-and-becn);
        translate-plp-control-word-de;
        unnumbered-address interface-name destination address destination-profile profile-name;
        vlan-id number;
        vlan-id-list [number number-number];
        address address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            destination address;
            destination-profile name;
            eui-64;
            primary-only;
            multipoint-destination address {
                dlci dlci-identifier;
                epd-threshold cells <plp1 cells>;
                inverse-arp;
                oam-liveness {
                    up-count cells;
                    down-count cells;
                }
                oam-period (disable | seconds);
                shaping {
                    (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length
```

```
peak rate sustained rate);
                    queue-length number;
                }
                vci vpi-identifier.vci-identifier;
            }
            preferred;
            primary;
            (vrrp-group | vrrp-inet6-group) group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                authentication-type authentication;
                authentication-key key;
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bits-per-second priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/prefix-length routing-instance instance-name priority-
cost cost;
                }
                virtual-address [addresses];
                virtual-link-local-address ipv6-address;
                vrrp-inherit-from {
                    active-interface interface-name;
                    active-group group-number;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name],
[edit interfaces interface-set interface-set-name interface interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 1,073,741,823 for demux, PPPoE, and pseudowire static interfaces. 0 through 16,385 for all other static interface types.

`etree-ac-role (leaf | root)`—To configure an interface as either leaf or root.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Range increased for static pseudowire interfaces to 1,073,741,823 in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- Configuring Logical Interface Properties
- Junos OS Services Interfaces Library for Routing Devices

unit (Dynamic Profiles Standard Interface)

IN THIS SECTION

- Syntax | 2479
- Hierarchy Level | 2482
- Description | 2482
- Options | 2482
- Required Privilege Level | 2483
- Release Information | 2483

Syntax

```
unit logical-unit-number {
    actual-transit-statistics;
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;

            }
            dynamic-profile profile-name;
        }
    }
    dial-options {
        ipsec-interface-id name;
        l2tp-interface-id name;
        (shared | dedicated);
    }
    encapsulation  (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-
mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux |
ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp |
```

```
ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-
relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-
over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
    family family {
        address address;
        demux-destination,
        filter {
            adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
            }
            input filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
            output filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
        }
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
            fail-filter filter-name;
            mode loose;
        }
        service {
             input {
                 service-set service-set-name {
                    service-filter filter-name;
                }
                post-service-filter filter-name;
            }
            input-vlan-map {
                inner-tag-protocol-id tpid;
                inner-vlan-id number;
                (push | swap);
                tag-protocol-id tpid;
                vlan-id number;
            }
```

```
            output {
                 service-set service-set-name {
                    service-filter filter-name;
                }
            }
            output-vlan-map {
                inner-tag-protocol-id tpid;
                inner-vlan-id number;
                (pop | swap);
                tag-protocol-id tpid;
                vlan-id number;
            }
        }
        service-name-table table-name
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
        input filter-name {
            shared-name  filter-shared-name;
        }
        output filter-name {
            shared-name  filter-shared-name;
        }
    }
    host-prefix-only;
    keepalives {
        interval seconds;
        }
    ppp-options {
        aaa-options aaa-options-name;
        authentication [ authentication-protocols ];
        chap {
            challenge-length minimum minimum-length maximum maximum-length;
            local-name name;
        }
        ignore-magic-number-mismatch;
        initiate-ncp (dual-stack-passive | ipv6 | ip)
        ipcp-suggest-dns-option;
        mru size;
        mtu (size | use-lower-layer);
        on-demand-ip-address;
        pap;
```

```
        peer-ip-address-optional;
        local-authentication {
            password password;
            username-include {
                circuit-id;
                delimiter character;
                domain-name name;
                mac-address;
                remote-id;
            }
        }
    }
    service {
        pcef  pcef-profile-name {
            activate rule-name | activate-all;
        }
    }
    targeted-options {
        backup backup;
        group group;
        primary primary;
        weight ($junos-interface-target-weight | weight-value);
    }
    vlan-id number;
    vlan-tags outer  [tpid].vlan-id [inner [tpid].vlan-id];
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following predefined variables:

- `$junos-underlying-interface-unit`—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- `$junos-interface-unit`—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Agent Circuit Identifier-Based Dynamic VLANs Overview

update-interval

IN THIS SECTION

- Syntax | 2484
- Hierarchy Level | 2484
- Description | 2484
- Default | 2484
- Options | 2485
- Required Privilege Level | 2485

- Release Information | 2485

Syntax

```
update-interval minutes;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Enable interim accounting updates and configure the amount of time that the router or switch waits before sending a new accounting update.

Interim accounting updates are included in the exchange of messages between the client and the accounting server. In RADIUS accounting, the client is the network access server (NAS), which can be the router or switch. The NAS sends Accounting-Request messages to the server, which acknowledges receipt of the requests with Accounting-Response messages. Interim accounting updates are sent in Accounting-Request packets with the Acct-Status-Type attribute set to Interim-Update.

When a user is authenticated, the authentication server issues an Access-Accept message in response to a successful Access-Request message. The interval between interim updates can be configured directly on the server using the Acct-Interim-Interval attribute of the Access-Accept message. However, if the update interval is configured on the NAS using `update-interval`, then the system prefers the attributes returned by RADIUS and overrides the locally configured values.

NOTE: All information in an interim update message is cumulative from the beginning of the session, not from the last interim update message.

Default

No interim updates are sent from the client to the accounting server.

Options

minutes—Amount of time between updates, in minutes. All values are rounded 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.

- **Range:** 10 through 1440 minutes

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 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access | 476

ANCP Agent and AAA | 1311

update-interval (Service Accounting)

IN THIS SECTION

- Syntax | 2486
- Hierarchy Level | 2486
- Description | 2486
- Default | 2486
- Options | 2486
- Required Privilege Level | 2486
- Release Information | 2487

Syntax

```
update-interval minutes;
```

Hierarchy Level

```
[edit access profile profile-name service accounting]
```

Description

Enable service interim accounting updates and configure the amount of time that the router waits before sending a new service accounting update.

NOTE: An update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26-140), takes precedence over the value configured with this statement. In turn, the value configured with this statement takes precedence over the *interval* value configured at the `[edit accounting-options flat-file-profile profile-name]` hierarchy level.

Default

No updates

Options

minutes—Amount of time between updates, in 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.

- **Range:** 10 through 1440 minutes

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 14.2R1.

RELATED DOCUMENTATION

- Configuring Service Accounting | 586
- Configuring Per-Subscriber Session Accounting | 573
- Processing Cisco VSAs in RADIUS Messages for Service Provisioning | 583
- *Configuring Service Accounting in Local Flat Files*

update-response-timeout (PCRF Partition)

IN THIS SECTION

- Syntax | 2487
- Hierarchy Level | 2487
- Description | 2488
- Options | 2488
- Required Privilege Level | 2488
- Release Information | 2488

Syntax

```
update-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the amount of time in seconds before a Policy and Charging Rule Function (PCRF) partition stops attempting to send an updated rule report response using a CCR-GX-U message. Use this statement if you added or deleted rules in either the subscriber login using a CCA-GX-I message, or the reauthorization request using a RAR-GX-U message, and configured rule reporting,

The last rule report may be preempted by the subscriber logging out too soon. Use this statement to control the amount of time to allow the system to send the last rule report.

Options

seconds Number of seconds to wait before a PCRF partition stops attempting to send an updated rule report response using a CCR-GX-U message.

- **Default:** 300
- **Range:** 0 through 86,400 seconds (24 hours)

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 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition | **1455**

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**

Understanding Gx Interactions Between the Router and the PCRF | **1418**

Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

Understanding Upstream and Downstream Messages for the PCRF | **1444**

upstream-rate (Traffic Shaping)

IN THIS SECTION

- Syntax | 2489
- Hierarchy Level | 2489
- Description | 2489
- Options | 2490
- Required Privilege Level | 2490
- Release Information | 2490

Syntax

```
upstream-rate rate;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit advisory-options],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name advisory-options],
[edit interfaces demux0 unit logical-unit-number advisory-options],
[edit interfaces interface-name logical-unit-number advisory-options]
```

Description

Specify a recommended shaping rate to be applied to upstream traffic on an interface.

For ANCP interfaces, this configured rate is used as the default value for the Juniper VSA Upstream-Calculated-Qos-Rate (26-142) when the router has not received and processed the attributes from the access node.

For L2TP, the rate is configured on an underlying PPPoE logical interface for a subscriber on an MX Series router acting as a LAC. When the subscriber is tunneled, this rate, referred to as speed for L2TP, is sent to the LNS in the ICCN message as AVP 38.

Options

rate—Traffic rate in bits per second.

- **Range:** 1000 through 4,294,967,295 bits per second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces | **1304**

Configuring the ANCP Agent | **1253**

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

use-interface-description

IN THIS SECTION

- Syntax | **2491**
- Hierarchy Level | **2491**
- Description | **2491**
- Options | **2492**
- Required Privilege Level | **2492**
- Release Information | **2492**

Syntax

```
use-interface-description (logical | device);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Use the textual interface description instead of the interface identifier in the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or in the DHCPv6 option 18 (Relay Agent Interface ID) or option 37 (Relay Agent Remote ID) information in DHCP packets that the DHCP relay agent sends to a DHCP server.

NOTE: For integrated routing and bridging (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.

The textual description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. If you specify that the textual description be used and no description is configured for the interface, DHCP relay defaults to using the Layer 2 interface name. When you use the interface

description rather than the interface name, the interface description has to be specified under interface unit ("set interfaces ge-0/0/0 unit 0 description "client"). If you do not do this, then the interface name is used.

In the case of integrated routing and bridging (IRB) interfaces, the textual description of the Layer 2 interface is used instead of the IRB interface. If there is no description configured, the Layer 2 logical interface name is used. To include the IRB interface description instead of the Layer 2 interface description, configure the `use-interface-description` and the `no-vlan-interface-name` statements. If no description is configured for the IRB interface, DHCP relay defaults to using the IRB interface name.

NOTE: The `use-interface-description` statement is mutually exclusive with the `use-vlan-id` statement.

If you specify the textual interface description, rather than accepting the default syntax, the identification is for packets returned from the server, and only for instances where that identification would be required by the DHCP relay, such as a stateless pass-through.

NOTE: By default, DHCP relay accepts a maximum of 253 ASCII characters. If the textual interface description exceeds 253 characters, DHCP relay drops the packet, which results in the DHCP client failing to bind.

Options

`logical`—Use the textual description that is configured for the logical interface.

`device`—Use the textual description that is configured for the device interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... relay-agent-remote-id]` and `[edit ... remote-id]` hierarchy levels introduced in Junos OS Release 14.1.

Support at the `[edit vlans vlan-name dhcp-security dhcpv6-options option-18]` and `[edit vlans vlan-name dhcp-security dhcpv6-options option-37]` hierarchy levels introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

- Using DHCP Relay Agent Option 82 Information | 751
- Using DHCP Relay Agent Option 82 Information | 751
- DHCPv6 Relay Agent | 916

username-include (Demux)

IN THIS SECTION

- Syntax | 2493
- Hierarchy Level | 2494
- Description | 2494
- Options | 2494
- Required Privilege Level | 2494
- Release Information | 2495

Syntax

```
username-include {
    auth-server-realm realm-string;
    delimiter delimiter-character;
    domain-name domain-name;
    interface-name;
    source-address;
    user-prefix user-prefix-string;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges
authentication]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges
authentication]
```

Description

Specify the information included in the username created for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Options

auth-server-realm *auth-server-realm-string*	Specify the authentication server name string used for the demultiplexing (demux) interface options.
delimiter *delimiter-character*	Specify the character used as the delimiter between the concatenated components of the username. You cannot use the semicolon (;) as a delimiter. The delimiter defines the boundary between the part of the original username that is kept and the part that is discarded.
domain-name *domain-name-string*	Specify the domain name formatted string that is concatenated with the username during the subscriber authentication process.
interface-name	Append the interface name for the demultiplexing (demux) interface options to the username string used for authentication.
source-address	Set the source address for every user name for the demultiplexing (demux) interface options.
user-prefix *user-prefix-string*	Specify the user prefix that is concatenated with the username for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview | 1110

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles | 1111

username-include (DHCP Local Server)

IN THIS SECTION

- Syntax | 2495
- Hierarchy Level | 2496
- Description | 2496
- Options | 2497
- Required Privilege Level | 2497
- Release Information | 2497

Syntax

```
username-include {
    circuit-type;
    client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interfaces-description (device-interface | logical-interface);
    interface-name ;
    logical-system-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
```

```
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...],
[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server group group-name authentication]
[edit system services dhcp-local-server dhcpv6 authentication],
[edit system services dhcp-local-server dhcpv6 group group-name authentication],
[edit system services dhcp-local-server dual-stack-group group-name authentication]
```

Description

Configure the username that the router or switch passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The following statements are not supported in the DHCPv6 hierarchy levels:

- **option-60**
- **option-82**

The following statements are supported in the DHCPv6 hierarchy levels only:

- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

Options

vlan-tags Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

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.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Specifying Authentication Support | 833

Creating Unique Usernames for DHCP Clients | 834

username-include (DHCP Relay Agent)

IN THIS SECTION

Syntax | 2498

Hierarchy Level | 2498

- Description | 2499
- Options | 2499
- Required Privilege Level | 2499
- Release Information | 2500

Syntax

```
username-include {
    circuit-type;
    client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (device-interface | logical-interface);
    interface-name;
    logical-system-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication],
[edit forwarding-options dhcp-relay group group-name authentication],
[edit forwarding-options dhcp-relay dhcpv6 authentication],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication],
[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 ...]
```

Description

Configure the username that the router (or switch) passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS. Use the statement at the `[edit...dhcpv6]` hierarchy levels to configure DHCPv6 support.

The following statements are not supported in the DHCPv6 hierarchy levels:

- **option-60**
- **option-82**

The following statements are supported in the DHCPv6 hierarchy levels only:

- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

Options

vlan-tags Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

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.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients | 834

Specifying Authentication Support | 833

user-name-include (OCS Partition)

IN THIS SECTION

- Syntax | 2500
- Hierarchy Level | 2501
- Description | 2501
- Options | 2501
- Required Privilege Level | 2502
- Release Information | 2502

Syntax

```
user-name-include{
    base-interface-name;
    delimiter delimiter-character;
    domain-name my-domain;
    interface-name;
    mac-address;
```

```
    nas-port-id;
    origin-host;
    origin-realm;
    user-name;
    user-prefix pref;
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the username identification to be used in an Online Charging Function (OCS) partition. You can configure a data string to include the following options.

Options

base-interface-name Use the physical or underlying interface name. If both the underlying interface and the client application perform authentication, authorization, and provisioning, the identification attributes in the server requests enable the AAA or PCRF server to make an association between the two entities.

delimiter *delimiter-character* Use the specified character as the delimiter between the concatenated components of the `user-name-include` statement. You cannot use the semicolon (;) as a delimiter.

- **Default:** @

domain-name *my-domain* Use the specified domain name that is concatenated with the `user-name-include` statement during the subscriber identification process.

interface-name Use the interface name that is concatenated with the `user-name-include` statement during the subscriber identification process; for example, demux0.

mac-address Use the client hardware address (chaddr) from the incoming packet that is concatenated with the `user-name-include` statement during the subscriber identification process.

nas-port-id Use the NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface that subscriber management uses to identify users that is concatenated with the `user-`

name-include statement during the subscriber identification process. By default, the NAS-Port-ID includes the interface-description value that describes the physical interface.

origin-host Use the name of the host that originates the Diameter message that is concatenated with the user-name-include statement during the subscriber identification process. Supplied as the value of Origin-Host AVP for all messages sent by the Diameter master instance.

origin-realm Use the realm of the host that originates the Diameter message that is concatenated with the user-name-include statement during the subscriber identification process. Supplied as the value of Origin-Realm AVP for all messages sent by the Diameter master instance.

user-name (Included in the user-name-include statement by default) Use the username that is concatenated with the user-name-include statement during the subscriber identification process.

user-prefix *prefix* Use the specified user prefix that is concatenated with the user-name-include statement during the subscriber identification process.

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 16.2.

RELATED DOCUMENTATION

- Configuring the OCS Partition | **1449**
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting | **1410**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

username-include (Static Subscribers)

IN THIS SECTION

- Syntax | 2503
- Hierarchy Level | 2503
- Description | 2504
- Options | 2504
- Required Privilege Level | 2504
- Release Information | 2504

Syntax

```
username-include {
    delimiter delimiter-character;
    domain-name domain-name;
    interface;
    logical-system-name;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication],
[edit logical-systems logical-system-name system services static-subscribers  authentication],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication],
[edit routing-instances routing-instances-name system services static-subscribers
authentication],
[edit routing-instances routing-instances-name system services static-subscribers group group-
```

```
name  authentication],
[edit system services static-subscribers authentication],
[edit system services static-subscribers group group-name authentication]
```

Description

Specify the information included in the username created for all static subscribers or for static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Options

delimiter *delimiter-character*

Specify the delimiter character that separates the different fields configured for the username. You can specify only a single character as the delimiter.

- **Default:** The default delimiter is a period (.).

vlan-tags

Include the VLAN tags (VLAN IDs) for the interface. For single-tagged VLANs, the field is encoded as *outer-tag*. For stacked VLANs, the field is encoded as *outer-tag-inner-tag*. For IP demux interfaces configured for static subscribers, the VLAN tags configured on the underlying interface are used.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`delimiter` and `vlan-tags` options added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview | 1487

Configuring the Static Subscriber Global Username | 1492

Configuring the Static Subscriber Group Username | 1497

use-option-82

IN THIS SECTION

- Syntax | 2505
- Hierarchy Level | 2505
- Description | 2506
- Options | 2506
- Required Privilege Level | 2506
- Release Information | 2506

Syntax

```
use-option-82 <strict>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id |relay-agent-remote-id)]
```

Description

For the `relay-agent-interface-id` statement, specify that in dual-stack environments, the DHCPv6 relay agent uses an Interface-ID option (option 18) that is based on the DHCPv4 relay agent information option (option 82). When you include this statement, the DHCPv6 relay agent checks for the option 82 Agent Circuit-ID suboption (suboption 1) and inserts it into the outgoing RELAY-FORW message.

For the `relay-agent-remote-id` statement, specify that in dual-stack environments, the DHCPv6 relay agent uses a Remote-ID option (option 37) that is based on the DHCPv4 relay agent information option (option 82). When you include this statement, the DHCPv6 relay agent checks for the option 82 Remote-ID suboption (suboption 2) and inserts it into the outgoing RELAY-FORW message.

Options

strict (Optional) The router drops Solicit messages that do not include a DHCPv4 Remote-ID (option 82 suboption 2). If you do not specify the `strict` keyword, the router sends the RELAY-FORW message without adding option 37.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support at the `[...relay-agent-remote-id]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information | **751**

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets | **919**

use-primary (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2507
- Hierarchy Level | 2507
- Description | 2508
- Options | 2508
- Required Privilege Level | 2508
- Release Information | 2508

Syntax

```
use-primary primary-profile-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dhcpv6 group  group-name dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name dynamic-profile profile-
name],
[edit forwarding-options dhcp-relay group  group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 dynamic-profile
profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dynamic-profile profile-
name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name dynamic-
profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dynamic-profile profile-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 dynamic-
profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dynamic-profile
profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group  group-
name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group  group-name
dynamic-profile profile-name]
```

Description

Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber logs in. Subsequent subscribers are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber logs out, the next subscriber that logs in is assigned the primary dynamic profile.

Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Options

primary-profile-name—Name of the dynamic profile to configure as the primary dynamic profile

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *dual-stack-group-name*`]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

use-underlying-interface-mac

IN THIS SECTION

- Syntax | 2509
- Hierarchy Level | 2509
- Description | 2509
- Required Privilege Level | 2510
- Release Information | 2510

Syntax

```
use-underlying-interface-mac
```

Hierarchy Level

```
[edit system demux-options]
```

Description

Configure the software to use the same link-local address for both interfaces when you are using Router Advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces. In this case, the link-local address for the underlying interface should be based the MAC address of the underlying interface.

This statement causes the system to assign an address using the 64-bit Extended Unique Identifier (EUI-64) as described in RFC 2373.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces | 1052

use-vlan-id

IN THIS SECTION

- Syntax | 2510
- For Platforms with Enhanced Layer 2 Software (ELS) | 2511
- For MX Series Platforms | 2511
- Description | 2511
- Required Privilege Level | 2511
- Release Information | 2511

Syntax

```
use-vlan-id;
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit forwarding-options helpers bootp dhcp-option82-circuit-id]
[edit forwarding-options helpers bootp interface  interface-name dhcp-option82-circuit-id]
```

For MX Series Platforms

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82 circuit-id]
```

Description

Use the VLAN ID rather than the VLAN name (the default) in the circuit ID or remote ID value in the DHCP option 82 information.

NOTE: The `use-vlan-id` statement is mutually exclusive with the `use-interface-description` and `no-vlan-interface-name` statements.

The `use-vlan-id` statement only applies to interfaces in a bridge domain. The format of the Agent Circuit ID or Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces is as follows:

```
(fe | ge)-fpc/pic/port.subunit:svlan_id-vlan_id
```

NOTE: The *subunit* is required and used to differentiate the interface for remote systems, and *svlan_id-vlan_id* represents the VLANs associated with the bridge domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

NOTE: The EX Series switches that support the `use-vlan-id` statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

- Example: Setting Up DHCP Option 82 Using the Same VLAN
- Example: Setting Up DHCP Option 82
- http://tools.ietf.org/html/rfc3046.

use-vlan-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **2512**
- Hierarchy Level | **2513**
- Description | **2513**
- Required Privilege Level | **2513**
- Release Information | **2513**

Syntax

```
use-vlan-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dual-stack-group dual-
stack-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dual-stack-group dual-stack-group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dual-stack-group ]
```

Description

Specify that the VLAN is identified by the VLAN ID, rather than the VLAN name, when you configure the DHCPv6 relay agent to include either of the following in packets it sends to a DHCPv6 server:

- DHCPv6 Interface-ID (option 18)
- DHCPv6 Remote-ID (option 37)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent | 916

Extended DHCP Relay Agent Overview | 695

user-prefix (DHCP Local Server)

IN THIS SECTION

- Syntax | 2514
- Hierarchy Level | 2514
- Description | 2515
- Options | 2515
- Required Privilege Level | 2515
- Release Information | 2516

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication ],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

user-prefix-string—User prefix string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

user-prefix (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2516
- Hierarchy Level | 2516
- Description | 2517
- Options | 2517
- Required Privilege Level | 2517
- Release Information | 2518

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-grouop dual-stack-group-name authentication
username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
```

include],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include]

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

user-prefix-string—User prefix string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCP with External Authentication Server

user-prefix (Static Subscribers)

IN THIS SECTION

- Syntax | **2518**
- Hierarchy Level | **2518**
- Description | **2519**
- Options | **2519**
- Required Privilege Level | **2519**
- Release Information | **2519**

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
```

```
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that a string is included as the beginning of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Options

user-prefix-string—String that begins the username. The string can include the following characters: a through z, A through Z, 0 through 9, "-", or ".".

Required Privilege Level

system—To view this statement in the configuration.

system-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 | 1487
- Configuring the Static Subscriber Global Username | 1492
- Configuring the Static Subscriber Group Username | 1497

valid-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2520
- Hierarchy Level | 2520
- Description | 2520
- Options | 2520
- Required Privilege Level | 2520
- Release Information | 2521

Syntax

```
valid-lifetime seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-name
prefix prefix]
```

Description

Specify how long the prefix remains valid for onlink determination.

Options

seconds—Valid lifetime, in seconds. If you set the valid lifetime to `0xffffffff`, the lifetime is infinite.

- **Default:** 2,592,000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA | 939

Dynamic Router Advertisement Configuration Overview | 942

vdsl-bytes

IN THIS SECTION

- Syntax | 2521
- Hierarchy Level | 2521
- Description | 2522
- Options | 2522
- Required Privilege Level | 2522
- Release Information | 2522

Syntax

```
vdsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for a VDSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

vdsl-overhead-adjust

IN THIS SECTION

- Syntax | 2523
- Hierarchy Level | 2523
- Description | 2523
- Options | 2524
- Required Privilege Level | 2524
- Release Information | 2524

Syntax

```
vdsl-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for a VDSL access line reported in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

vdsl2-bytes

IN THIS SECTION

- Syntax | 2525
- Hierarchy Level | 2525
- Description | 2525
- Options | 2526
- Required Privilege Level | 2526
- Release Information | 2526

Syntax

```
vdsl2-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for a VDSL2 access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl2-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS | 1299

Traffic Rate Reporting and Adjustment by the ANCP Agent | 1293

Configuring the ANCP Agent | 1253

vdsl2-overhead-adjust

IN THIS SECTION

Syntax | 2527

- Hierarchy Level | 2527
- Description | 2527
- Options | 2527
- Required Privilege Level | 2528
- Release Information | 2528

Syntax

```
vdsl2-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for a VDSL2 access line received in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl2-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

- Configuring the ANCP Agent to Report Traffic Rates to CoS | **1299**
- Traffic Rate Reporting and Adjustment by the ANCP Agent | **1293**
- Configuring the ANCP Agent | **1253**

vendor-specific (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **2529**
- Hierarchy Level | **2529**
- Description | **2529**
- Required Privilege Level | **2529**
- Release Information | **2529**

Syntax

```
vendor-specific {
    host-name;
    location;
}
```

Hierarchy Level

```
[edit forwarding-optionsdhcp-relay relay-option-82],
[edit forwarding-optionsdhcp-relay group  group-name relay-option-82]
```

Description

Add vendor-specific information to the option-82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The vendor-specific information can be a hostname, a location (such as a unique connection identifier), or both. The hostname can be a string of characters such as **Juniper-AB-1**. The location should be specified as interface, VLAN ID, and if applicable, stacked VLAN ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243. The enterprise ID is 2636. The hostname is option-data 1, and the location is option-data 2. The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query a single entity to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

violation-action (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2530
- Hierarchy Level | 2530
- Description | 2530
- Options | 2531
- Required Privilege Level | 2531
- Release Information | 2531

Syntax

```
violation-action action;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 lease-time-validation],
[edit forwarding-options dhcp-relay group group-name lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 group group-name lease-time-validation],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server lease-time-validation],
[edit system services dhcp-local-server dhcpv6 lease-time-validation],
[edit system services dhcp-local-server group group-name lease-time-validation],
[edit system services dhcp-local-server dhcpv6 group group-name lease-time-validation]
```

Description

Configure the action that the router performs when a DHCP lease-time violation occurs. The violation occurs when a third-party DHCP server or address-assignment pool offers a DHCP lease time that is less than the threshold specified by the `lease-time-threshold` statement.

Options

action Action taken by the router when a lease-time violation occurs.

- `drop`—(Optional) For DHCPv4 and DHCPv6 relay agent, the third-party lease is dropped and the client binding fails.
- `override-lease`—(Optional) For DHCPv4 and DHCPv6 local server, the third-party lease is overridden with the value specified by the `lease-time-threshold` statement and binds the client using the new value.
- `strict`—(Optional) For DHCPv4 and DHCPv6 local server, DHCP ignores the third-party lease and the client binding fails.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold | 783

vlan-ranges (RADIUS Options)

IN THIS SECTION

Syntax | 2532

- Hierarchy Level | 2532
- Description | 2532
- Options | 2532
- Required Privilege Level | 2533
- Release Information | 2533

Syntax

```
vlan-ranges (any | low-tag-high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the VLAN range of subscribers to which the named NAS-Port options definition applies.

NOTE: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

any Entire VLAN range representing all VLAN IDs.

low-tag VLAN ID tag representing the lower limit of the VLAN range.

- **Range:** 1 through 4094

high-tag VLAN ID tag representing the upper limit of the VLAN range.

- **Range:** 1 through 4094

NOTE: To specify a single VLAN ID, set *low-tag* and *high-tag* to the same value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **527**

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN | **526**

vrf-name (Duplicate Accounting)

IN THIS SECTION

- Syntax | **2534**
- Hierarchy Level | **2534**
- Description | **2534**
- Options | **2534**
- Required Privilege Level | **2534**
- Release Information | **2534**

Syntax

```
vrf-name vrf-name;
```

Hierarchy Level

```
[edit access profile profile-name accounting duplication-vrf]
```

Description

Specify a nondefault VRF (LS:RI combination) to which duplicate accounting information is sent. Up to five access profiles can be defined in this VRF; the profiles point to the RADIUS accounting servers that receive the accounting information.

Options

vrf-name Name of a nondefault VRF to receive duplicate accounting reports.

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 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

Understanding RADIUS Accounting Duplicate Reporting | 578

Configuring Authentication and Accounting Parameters for Subscriber Access | 550

wait-for-acct-on-ack (Access Profile)

IN THIS SECTION

- Syntax | 2535
- Hierarchy Level | 2535
- Description | 2535
- Required Privilege Level | 2536
- Release Information | 2536

Syntax

```
wait-for-acct-on-ack;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router's authd process to wait for an Acct-On-Ack response message from RADIUS before sending 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.

At subscriber session startup, the Junos OS authd process sends an Acct-On message to the RADIUS server and the new session starts authentication and accounting operations. However, in some service provider environments, upon receipt of the Acct-On message, the RADIUS server cleans up the previous session state and removes accounting statistics. In this scenario, the RADIUS server's cleanup operation can inadvertently delete the new session's authentication and accounting information, which might include customer billing information.

To ensure that the new session's authentication and accounting information is not deleted, you can include the `wait-for-acct-on-ack` statement to configure authd to wait for an Acct-On-Ack response

message from the RADIUS accounting server, so the RADIUS cleanup can finish before authd sends any new authentication and accounting updates.

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 12.3.

RELATED DOCUMENTATION

Configuring Per-Subscriber Session Accounting | 573

RADIUS Servers and Parameters for Subscriber Access | 476

CHAPTER 14

Operational Commands

IN THIS CHAPTER

- clear ancp neighbor | 2541
- clear ancp statistics | 2543
- clear ancp subscriber | 2547
- clear dhcp relay active-leasequery statistics | 2549
- clear dhcp relay binding | 2553
- clear dhcp relay lockout-entries | 2556
- clear dhcp relay statistics | 2559
- clear dhcp server active-leasequery statistics | 2562
- clear dhcp server binding | 2563
- clear dhcp server lockout-entries | 2568
- clear dhcp server statistics | 2570
- clear dhcpv6 relay active-leasequery statistics | 2573
- clear dhcpv6 relay binding | 2576
- clear dhcpv6 relay lockout-entries | 2580
- clear dhcpv6 relay statistics | 2583
- clear dhcpv6 server active-leasequery statistics | 2585
- clear dhcpv6 server binding | 2587
- clear dhcpv6 server lockout-entries | 2591
- clear dhcpv6 server statistics | 2593
- clear diameter function statistics | 2595
- clear diameter peer | 2597
- clear extensible-subscriber-services counters | 2598
- clear extensible-subscriber-services sessions | 2599
- clear ipv6 router-advertisement | 2601
- clear network-access aaa statistics | 2602
- clear network-access aaa subscriber | 2607

- clear network-access gx-plus replay | **2610**
- clear network-access gx-plus statistics | **2611**
- clear network-access ocs statistics | **2613**
- clear network-access pcrf | **2614**
- clear system subscriber-management statistics | **2616**
- request ancp oam interface | **2617**
- request ancp oam neighbor | **2619**
- request ancp oam port-down | **2621**
- request ancp oam port-up | **2623**
- request dhcp relay bulk-leasequery | **2625**
- request dhcp relay leasequery | **2628**
- request dhcp server reconfigure | **2630**
- request dhcpv6 server reconfigure | **2632**
- request dhcpv6 relay bulk-leasequery | **2634**
- request dhcpv6 relay leasequery | **2637**
- request network-access aaa accounting | **2639**
- request network-access aaa replay pending-accounting-stops | **2641**
- request network-access aaa subscriber modify session-id | **2642**
- request network-access aaa subscriber set session-id | **2644**
- request services extensible-subscriber-services reload-dictionary | **2646**
- request services static-subscribers login group | **2648**
- request services static-subscribers logout group | **2649**
- request services static-subscribers login interface | **2651**
- request services static-subscribers logout interface | **2652**
- request services subscribers | **2654**
- request services subscribers clear | **2655**
- request system reboot (Junos OS) | **2657**
- restart extensible-subscriber-services | **2667**
- show accounting pending-accounting-stops | **2669**
- show ancp cos | **2674**
- show ancp neighbor | **2681**
- show ancp statistics | **2693**

- show ancp subscriber | 2699
- show ancp summary | 2709
- show ancp summary neighbor | 2712
- show ancp summary subscriber | 2715
- show class-of-service interface | 2717
- show class-of-service interface-set | 2762
- show class-of-service scheduler-map | 2765
- show class-of-service traffic-control-profile | 2770
- show database-replication statistics | 2776
- show database-replication summary | 2778
- show dhcp relay active-leasequery | 2782
- show dhcp relay binding | 2789
- show dhcp relay lockout-entries | 2796
- show dhcp relay statistics | 2800
- show dhcp server active-leasequery | 2807
- show dhcp server active-leasequery statistics | 2809
- show dhcp server binding | 2811
- show dhcp server lockout-entries | 2821
- show dhcp server statistics | 2825
- show dhcpv6 relay active-leasequery | 2831
- show dhcpv6 relay binding | 2838
- show dhcpv6 relay lockout-entries | 2850
- show dhcpv6 relay statistics | 2854
- show dhcpv6 server active-leasequery | 2860
- show dhcpv6 server active-leasequery statistics | 2862
- show dhcpv6 server binding | 2865
- show dhcpv6 server lockout-entries | 2874
- show dhcpv6 server statistics | 2877
- show diameter | 2883
- show diameter function | 2892
- show diameter function statistics | 2898
- show diameter instance | 2903

- show diameter network-element | **2905**
- show diameter network-element map | **2910**
- show diameter peer | **2914**
- show diameter peer map | **2920**
- show diameter peer statistics | **2924**
- show diameter route | **2929**
- show dynamic-profile session | **2932**
- show ipv6 router-advertisement | **2939**
- show network-access aaa accounting | **2946**
- show network-access aaa radius-servers | **2948**
- show network-access aaa statistics | **2958**
- show network-access aaa statistics authentication | **2974**
- show network-access aaa statistics pending-accounting-stops | **2979**
- show network-access aaa statistics preauthentication | **2980**
- show network-access aaa statistics re-authentication | **2983**
- show network-access aaa subscribers | **2985**
- show network-access aaa subscribers session-id | **2992**
- show network-access aaa terminate-code | **3003**
- show network-access address-assignment pool | **3010**
- show network-access domain-map | **3012**
- show network-access gx-plus | **3014**
- show network-access nasreq statistics | **3032**
- show network-access ocs | **3037**
- show network-access pcrf | **3040**
- show network-access requests statistics | **3045**
- show network-access s6a | **3048**
- show ppp address-pool | **3051**
- show static-subscribers sessions | **3053**
- show subscribers | **3057**
- show subscribers summary | **3110**
- show system subscriber-management redundancy-state dhcp active-leasequery interface | **3121**
- show system subscriber-management redundancy-state interface | **3124**

- show system subscriber-management route | 3127
- show system subscriber-management statistics | 3134
- show system subscriber-management summary | 3147
- test aaa authd-lite user | 3151
- test aaa dhcp user | 3156
- test aaa ppp user | 3164

clear ancp neighbor

IN THIS SECTION

- Syntax | 2541
- Description | 2541
- Options | 2542
- Required Privilege Level | 2542
- Output Fields | 2542
- Sample Output | 2542
- Release Information | 2543

Syntax

```
clear ancp neighbor
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic

rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is *enabled*) and can reestablish adjacencies.

This command initiates logout of ANCP-triggered dynamic VLAN sessions on the physical interface associated with the specified neighbor; conventionally autosensed dynamic VLAN sessions and their associated logical interfaces are not affected.

Options

none	Clear all ANCP neighbors.
ip-address *ip-address*	(Optional) Clear the ANCP neighbor specified by the IP address.
system-name *mac-address*	(Optional) Clear the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor` command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

clear ancp neighbor

```
user@host> clear ancp neighbor
```

show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the `clear ancp neighbor` command was issued.

```
user@host> show ancp neighbor
```

```
  IP Address        MAC Address         State         Subscriber   Capabilities
                                                        Count
  203.0.113.102     00:00:5e:00:53:10  Established    5              Topo
  203.0.113.122     00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132     00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142     00:00:5e:00:53:14  Established    5              Topo


user@host> clear ancp neighbor ip-address 203.0.113.102

user@host> show ancp neighbor

  IP Address        MAC Address         State         Subscriber   Capabilities
                                                        Count
  203.0.113.122     00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132     00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142     00:00:5e:00:53:14  Established    5              Topo
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ancp neighbor | 2681

clear ancp statistics

IN THIS SECTION

- Syntax | 2544
- Description | 2544
- Options | 2544
- Required Privilege Level | 2544

- Output Fields | **2544**
- Sample Output | **2545**
- Release Information | **2546**

Syntax

```
clear ancp statistics
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear current statistics accumulated by the ANCP agent for all ANCP neighbors or the specified neighbor.

Options

none	Clear all ANCP statistics.
ip-address *ip-address*	(Optional) Clear statistics for the ANCP neighbor specified by the IP address.
system-name *mac-address*	(Optional) Clear statistics for the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor` command before and after clearing the ANCP neighbor statistics to verify the clear operation.

Sample Output

clear ancp statistics

```
user@host> clear ancp statistics
```

show ancp neighbor

The following sample output displays statistics for an ANCP neighbor before and after the `clear ancp statistics` command was issued.

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                             : 38
          TCP Port                            : 64959
          State                               : Established
          Subscriber Count                    : 7
          Capabilities                        : Topology Discovery
          System Instance                     : 11
          Peer Instance                       : 1
          Adjacency Timer (in 100ms)          : 50
          Peer Adjacency Timer (in 100ms)     : 100
          Partition Type                      : 0
          Partition Flag                      : 1
          Partition Identifier                : 0
          Dead Timer                          : 22
          Received Syn Count                  : 47
          Received Synack Count               : 48
          Received Rstack Count               : 2
          Received Ack Count                  : 12
          Received Port Up Count              : 8
          Received Port Down Count            : 2
          Received Other Count                : 0
          Sent Syn Count                      : 48
          Sent Synack Count                   : 47
          Sent Rstack Count                   : 1
          Sent Ack Count                      : 12
```

```
          Max Discovery Limit Exceed Count    : 0

user@host> clear ancp statistics ip-address 192.168.10.1

user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                           : 38
          TCP Port                          : 64959
          State                             : Established
          Subscriber Count                  : 7
          Capabilities                      : Topology Discovery
          System Instance                   : 11
          Peer Instance                     : 1
          Adjacency Timer (in 100ms)        : 50
          Peer Adjacency Timer (in 100ms)   : 100
          Partition Type                    : 0
          Partition Flag                    : 1
          Partition Identifier              : 0
          Dead Timer                        : 22
          Received Syn Count                : 0
          Received Synack Count             : 0
          Received Rstack Count             : 0
          Received Ack Count                : 0
          Received Port Up Count            : 0
          Received Port Down Count          : 0
          Received Other Count              : 0
          Sent Syn Count                    : 0
          Sent Synack Count                 : 0
          Sent Rstack Count                 : 0
          Sent Ack Count                    : 0
          Max Discovery Limit Exceed Count    : 0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show ancp neighbor | 2681

clear ancp subscriber

IN THIS SECTION

- Syntax | 2547
- Description | 2547
- Options | 2547
- Required Privilege Level | 2548
- Output Fields | 2548
- Sample Output | 2548
- Release Information | 2549

Syntax

```
clear ancp subscriber
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear the ANCP agent connection with all ANCP subscribers or with the specified ANCP subscriber. This command deletes information for the subscribers, causing the adjusted traffic rate to revert to the configured rate for the subscriber interface, but otherwise has no affect on ANCP neighbors.

Options

none	Clear all ANCP subscribers.

identifier *identifier-string*	(Optional) Clear the ANCP subscriber identified by the access loop ID.
ip-address *ip-address*	(Optional) Clear all ANCP subscribers on the neighbor specified by the IP address.
system-name *mac-address*	(Optional) Clear all ANCP subscribers on the neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp subscriber` command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

show ancp subscriber brief

```
user@host> show ancp subscriber brief
Loop Identifier          Type      Interface        Rate          Neighbor
                                                     Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102
  port-1-11                VDSL2    set-ge-10411       64       203.0.113.112
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.122
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.123
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
user@host> clear ancp subscriber identifier port-2-10

user@host> show ancp subscriber brief
Loop Identifier          Type      Interface        Rate          Neighbor
                                                     Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102
  port-1-11                VDSL2    set-ge-10411       64       203.0.113.112
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
```

clear ancp subscriber

```
user@host> clear ancp subscriber
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show ancp subscriber | 2699

clear dhcp relay active-leasequery statistics

IN THIS SECTION

- Syntax | 2549
- Description | 2550
- Options | 2550
- Required Privilege Level | 2550
- Output Fields | 2550
- Sample Output | 2550
- Sample Output | 2551
- Release Information | 2552

Syntax

```
clear dhcp relay active-leasequery statistics
(interface interface-name | peer ip-address)
```

Description

Clear active leasequery statistics for the specified access interface or for the specified DHCPv4 active leasequery peer relay agent. The statistics for bindings sent, received, installed and failed are cleared.

Options

interface *interface-name* (Optional) Clear all active leasequery statistics for a specific access interface.

peer *ip-address* (Optional) Clear all active leasequery statistics for a specific peer relay agent.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020

     Interface : ge-2/1/1.1020
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 400
     Bindings Received                     : 320
     Bindings Installed Successfully       : 310
     Bindings Failed to install            : 10
     Last Synchronization Time             : 2019-02-15 16:20:05 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
     Local Interface count                 : 1
     Remote Interface count                : 1
```

clear dhcp relay active-leasequery statistics (Interface)

```
user@host> clear dhcp relay active-leasequery statistics interface ge-2/1/1.1020
```

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020

     Interface : ge-2/1/1.1020
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 0
     Bindings Received                     : 0
     Bindings Installed Successfully       : 0
     Bindings Failed to install            : 0
     Last Synchronization Time             : 2019-02-15 16:21:58 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
     Local Interface count                 : 1
     Remote Interface count                : 1
```

Sample Output

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2

     peer : 192.0.2.2
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 225
     Bindings Received                     : 210
     Bindings Installed Successfully       : 210
     Bindings Failed to install            : 0
     Last Synchronization Time             : 2019-02-15 16:28:36 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
```

```
    Local Interface count                : 4
    Remote Interface count               : 4
```

clear dhcp relay active-leasequery statistics (Peer)

```
user@host> clear dhcp relay active-leasequery statistics
```

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2

    peer : 192.0.2.2
    Topology-Discover Configured         : Yes
    State                                : Done
    Bindings Sent                        : 0
    Bindings Received                    : 0
    Bindings Installed Successfully      : 0
    Bindings Failed to install           : 0
    Last Synchronization Time            : 2019-02-15 16:30:01 IST
    ALQ Transmit Buffer count            : ffff
    Max Leasequery Transmit Rate         : 60
    Local Interface count                : 4
    Remote Interface count               : 4
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show dhcp relay active-leasequery | 2782

clear dhcpv6 relay active-leasequery statistics | 2573

clear dhcp relay binding

IN THIS SECTION

- Syntax | 2553
- Description | 2553
- Options | 2553
- Required Privilege Level | 2554
- Output Fields | 2554
- Sample Output | 2554
- Release Information | 2556

Syntax

```
clear dhcp relay binding
<address>
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.

Options

address (Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.

- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all	(Optional) Clear the binding state for all DHCP clients.
dual-stack	(Optional) Clear the binding state for DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.
interface *interface-name*	(Optional) Clear the binding state for DHCP clients on the specified interface.
interfaces-vlan	(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Clear the binding state for DHCP clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See "show dhcp relay binding" on page 2789 for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the `clear dhcp relay binding` command is issued.

```
user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
198.51.100.32     00:00:5e:00:53:01  active   2007-02-08 16:41:17 EST
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 198.51.100.32

user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding dual-stack all

The following command clears all DHCP relay agent bindings for all DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

```
user@host> clear dhcp relay binding dual-stack all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

Release Information

Command introduced in Junos OS Release 8.3.

Options `all` and `interface` added in Junos OS Release 8.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Option `dual-stack` added in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

show dhcp relay binding | 2789

clear dhcp relay lockout-entries

IN THIS SECTION

Syntax | 2557

Description | 2557

- Options | 2557
- Required Privilege Level | 2557
- Output Fields | 2557
- Sample Output | 2558
- Release Information | 2558

Syntax

```
clear dhcp relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all	Clear all client entries from the lockout database.
index *index*	Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcp relay lockout-entries" on page 2796 for an explanation of the output fields.

Sample Output

clear dhcp relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    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     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for all entries.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp relay lockout-entries all

user@host> show dhcp relay lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcp relay lockout-entries | 2796

clear dhcp relay statistics

IN THIS SECTION

- Syntax | 2559
- Syntax | 2559
- Description | 2560
- Options | 2560
- Required Privilege Level | 2560
- Output Fields | 2560
- Sample Output | 2560
- Release Information | 2561

Syntax

```
clear dhcp relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Syntax

Syntax for EX Series switches:

```
show dhcp relay statistics
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCP relay bulk leasequery statistics.
leasequery	(Optional) Clear DHCP relay individual leasequery statistics.
logical-system *logical-system-name*	(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See `show dhcp relay statistics` for an explanation of output fields.

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the `clear dhcp relay statistics` command is issued.

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                     1
    Lease Time Violated       1

Messages received:
    BOOTREQUEST               116
```

```
    DHCPDECLINE                0
    DHCPDISCOVER               11
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                105

Messages sent:
    BOOTREPLY                  44
    DHCPOFFER                  11
    DHCPACK                    11
    DHCPNAK                    11


user@host> clear dhcp relay statistics

user@host> show dhcp relay statistics
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

Release Information

Command introduced in Junos OS Release 8.3.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show dhcp relay statistics | 2800

clear dhcp server active-leasequery statistics

IN THIS SECTION

- Syntax | 2562
- Description | 2562
- Required Privilege Level | 2562
- Sample Output | 2563
- Release Information | 2563

Syntax

```
clear dhcp server active-leasequery statistics
```

Description

Clear the active leasequery statistics of the DHCP local server.

Required Privilege Level

clear

Sample Output

clear dhcp server active-leasequery statistics

```
user@host> clear dhcp server active-leasequery statistics peer 192.0.2.0
```

```
    peer : 192.0.2.0
    Topology-Discover Configured          : Yes
    State                                 : Done
    Bindings Sent                         : 0
    Bindings Received                     : 0
    Bindings Installed Successfully       : 0
    Bindings Failed to install            : 0
    Last Synchronization Time             : 1970-01-01 05:30:00 IST
    ALQ Transmit Buffer count             : ffff
    Max Leasequery Transmit Rate          : 60
    Local Interface count                 : 1
    Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show dhcp server active-leasequery statistics | 2809

clear dhcp server binding

IN THIS SECTION

- Syntax | 2564
- Description | 2564

- Options | 2564
- Required Privilege Level | 2565
- Output Fields | 2566
- Sample Output | 2566
- Release Information | 2567

Syntax

```
clear dhcp server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on 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.

Options

address (Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.

- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCP clients.

interface ***interface-name*** (Optional) Clear the binding state for DHCP clients on the specified interface.

> **NOTE**: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

interfaces-vlan (Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard (Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name*** (Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance ***routing-instance-name*** (Optional) Clear the binding state for DHCP clients on the specified routing instance.

dual-stack (Optional) Remove either both arms or single arm of dual-stack.

> **NOTE**:
>
> - The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
> - When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

view

Output Fields

See "show dhcp server binding" on page 2811 for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the `clear dhcp server binding` command is issued.

```
user@host> show dhcp server binding

2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address       Hardware address   Type     Lease expires at
198.51.100.1      00:00:5e:00:53:01  active   2007-01-17 11:38:47 PST
198.51.100.3      00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST


user@host> clear dhcp server binding 198.51.100.1

user@host> show dhcp server binding

1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address       Hardware address   Type     Lease expires at
198.51.100.3      00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST
```

clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

clear dhcp server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcp server binding dual-stack all
```

Release Information

Command introduced in Junos OS Release 9.0.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with `dual-stack` statement in Junos OS Release 17.3.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

show dhcp server binding | 2811

clear dhcp server lockout-entries

IN THIS SECTION

- Syntax | 2568
- Description | 2568
- Options | 2568
- Required Privilege Level | 2569
- Output Fields | 2569
- Sample Output | 2569
- Release Information | 2570

Syntax

```
clear dhcp server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcp server lockout-entries" on page 2821 for an explanation of the output fields.

Sample Output

clear dhcp server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp server lockout-entries index 2

user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for all entries.

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
```

```
1     00:00:5E:00:53:00     Lockout     30          5200          2
2     00:00:5E:00:53:11     Grace      120          780           2
3     00:00:5E:00:53:22     Lockout    180          2300          1

user@host> clear dhcp server lockout-entries all

user@host> show dhcp server lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcp server lockout-entries | 2821

clear dhcp server statistics

IN THIS SECTION

- Syntax | 2570
- Description | 2571
- Options | 2571
- Required Privilege Level | 2571
- Output Fields | 2571
- Sample Output | 2571
- Release Information | 2572

Syntax

```
clear dhcp server statistics
<bulk-leasequery-connections>
```

```
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCPv4 local server bulk leasequery statistics.
logical-system *logical-system-name*	(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See "show dhcp server statistics" on page 2825 for an explanation of output fields.

Sample Output

clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the `clear dhcp server statistics` command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
    Total                     1
    Lease Time Violation      1
```

```
Messages received:
    BOOTREQUEST                89163
    DHCPDECLINE                0
    DHCPDISCOVER               8110
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                81053

Messages sent:
    BOOTREPLY                  32420
    DHCPOFFER                  8110
    DHCPACK                    8110
    DHCPNAK                    8100


user@host> clear dhcp server statistics

user@host> show dhcp server statistics
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

Release Information

Command introduced in Junos OS Release 9.0.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

clear dhcpv6 relay active-leasequery statistics

IN THIS SECTION

- Syntax | 2573
- Description | 2573
- Options | 2573
- Required Privilege Level | 2573
- Output Fields | 2574
- Sample Output | 2574
- Sample Output | 2575
- Release Information | 2576

Syntax

```
clear dhcpv6 relay active-leasequery statistics
(interface interface-name | peer ipv6-address)
```

Description

Clear active leasequery statistics for the specified access interface or for the specified DHCPv6 active leasequery peer relay agent. The statistics for bindings sent, received, installed and failed are cleared.

Options

interface *interface-name* (Optional) Clear all active leasequery statistics for a specific access interface.

peer *ip-address* (Optional) Clear all active leasequery statistics for a specific peer relay agent.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1

     Interface : ge-0/0/0.1
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 325
     Bindings Received                     : 150
     Bindings Installed Successfully       : 150
     Bindings Failed to install            : 0
     Last Synchronization Time             : 2019-02-28 14:11:32 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
     Local Interface count                 : 1
     Remote Interface count                : 1
```

clear dhcp relay active-leasequery statistics (Interface)

```
user@host> clear dhcp relay active-leasequery statistics interface ge-0/0/0.1
```

show dhcpv6 relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1

     Interface : ge-0/0/0.1
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 0
     Bindings Received                     : 0
     Bindings Installed Successfully       : 0
     Bindings Failed to install            : 0
```

```
     Last Synchronization Time              : 2019-02-28 14:12:10 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

Sample Output

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 420
     Bindings Received                      : 220
     Bindings Installed Successfully        : 220
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-03-22 09:15:15 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 4
     Remote Interface count                 : 4
```

clear dhcp relay active-leasequery statistics (Peer)

```
user@host> clear dhcp relay active-leasequery statistics
```

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 0
     Bindings Received                      : 0
```

```
Bindings Installed Successfully        : 0
Bindings Failed to install             : 0
Last Synchronization Time              : 2019-03-22 09:17:05 IST
ALQ Transmit Buffer count              : ffff
Max Leasequery Transmit Rate           : 60
Local Interface count                  : 4
Remote Interface count                 : 4
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show dhcpv6 relay active-leasequery | 2831

clear dhcp relay active-leasequery statistics | 2549

clear dhcpv6 relay binding

IN THIS SECTION

- Syntax | 2576
- Description | 2577
- Options | 2577
- Required Privilege Level | 2578
- Output Fields | 2578
- Sample Output | 2578
- Release Information | 2580

Syntax

```
clear dhcpv6 relay binding
<address>
```

```
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.

Options

address	(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries: • *CID*—The specified Client ID (CID). • *ipv6-prefix*—The specified IPv6 prefix. • *session-id*—The specified session ID.
all	(Optional) Clear the binding state for all DHCPv6 clients.
dual-stack	(Optional) Clear the binding state for DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.
interfaces-vlan	(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).
interface *interface-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
logical-system *logical-system-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See `show dhcpv6 relay binding` for an explanation of output fields.

Sample Output

clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the `clear dhcpv6 relay binding` command is issued.

```
user@host> show dhcpv6 relay binding

Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64   1           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64   2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64   3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64   4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64   5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64   6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:db8:3c4d:15::/64
user@host> show dhcpv6 relay binding

Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:16::/64   2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64   3           83720    BOUND    ge-1/0/0.0
```

```
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

clear dhcpv6 relay binding dual-stack all

The following command clears all DHCPv6 relay agent bindings for all DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.

```
user@host> clear dhcpv6 relay binding dual-stack all
```

clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 relay binding interface ae0
```

clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

Release Information

Command introduced in Junos OS Release 11.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Option `dual-stack` added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings | 896

show dhcpv6 relay binding | 2838

clear dhcpv6 relay lockout-entries

IN THIS SECTION

Syntax | 2581

- Description | 2581
- Options | 2581
- Required Privilege Level | 2581
- Output Fields | 2581
- Sample Output | 2582
- Release Information | 2582

Syntax

```
clear dhcpv6 relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all	Clear all client entries from the lockout database.
index *index*	Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcpv6 relay lockout-entries" on page 2850 for an explanation of the output fields.

Sample Output

clear dhcpv6 relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcpv6 relay lockout-entries index 2

user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcpv6 relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcpv6 relay lockout-entries all

user@host> show dhcpv6 relay lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcpv6 relay lockout-entries | 2850

clear dhcpv6 relay statistics

IN THIS SECTION

- Syntax | 2583
- Description | 2583
- Options | 2583
- Required Privilege Level | 2584
- Output Fields | 2584
- Sample Output | 2584
- Release Information | 2585

Syntax

```
clear dhcpv6 relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCPv6 relay bulk leasequery statistics.
leasequery	(Optional) Clear DHCPv6 relay individual leasequery statistics.

logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See `show dhcpv6 relay statistics` for an explanation of output fields.

Sample Output

clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the `clear dhcpv6 relay statistics` command is issued.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0
    Lease Time Violated         1

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
user@host> clear dhcpv6 relay statistics
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              0
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              0
    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
```

Release Information

Command introduced in Junos OS Release 11.4.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

clear dhcpv6 server active-leasequery statistics

IN THIS SECTION

Syntax | 2586

- Description | 2586
- Options | 2586
- Required Privilege Level | 2586
- Sample Output | 2587
- Release Information | 2587

Syntax

```
clear dhcpv6 server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Clear the active leasequery statistics of the DHCPv6 local server.

Options

peer *peer-address*	IP address of the peer DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the peer DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

clear

Sample Output

clear dhcpv6 server active-leasequery statistics peer ip-address

```
user@host> clear dhcpv6 server active-leasequery statistics peer 2001:db8:3::2
```

```
    peer : 2001:db8:3::2
    Topology-Discover Configured          : Yes
    State                                 : Done
    Bindings Sent                         : 0
    Bindings Received                     : 0
    Bindings Installed Successfully       : 0
    Bindings Failed to install            : 0
    Last Synchronization Time             : 1970-01-01 05:30:00 IST
    ALQ Transmit Buffer count             : ffff
    Max Leasequery Transmit Rate          : 60
    Local Interface count                 : 1
    Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show dhcpv6 server active-leasequery statistics | 2862

clear dhcpv6 server binding

IN THIS SECTION

- Syntax | 2588
- Description | 2588

- Options | 2588
- Required Privilege Level | 2589
- Output Fields | 2589
- Sample Output | 2589
- Release Information | 2591

Syntax

```
clear dhcpv6 server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.

Options

address (Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCPv6 clients.

interface *interface-name* (Optional) Clear the binding state for DHCPv6 clients on the specified interface.

interfaces-vlan	(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.
dual-stack	(Optional) Remove either both arms or single arm of dual-stack.

NOTE:

- The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
- When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcpv6 server binding dual-stack all
```

Release Information

Command introduced in Junos OS Release 9.6.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with `dual-stack` statement in Junos OS Release 17.3.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

show dhcpv6 server binding | 2865

clear dhcpv6 server lockout-entries

IN THIS SECTION

- Syntax | 2591
- Description | 2591
- Options | 2592
- Required Privilege Level | 2592
- Output Fields | 2592
- Sample Output | 2592
- Release Information | 2593

Syntax

```
clear dhcpv6 server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared.

The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcpv6 server lockout-entries" on page 2874 for an explanation of the output fields.

Sample Output

clear dhcpv6 server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 server 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 dhcpv6 server lockout-entries index 2

user@host> show dhcpv6 server 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
```

clear dhcpv6 server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 server 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 dhcpv6 server lockout-entries all

user@host> show dhcpv6 server lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcpv6 server lockout-entries | 2874

clear dhcpv6 server statistics

IN THIS SECTION

- Syntax | 2594
- Description | 2594
- Options | 2594
- Required Privilege Level | 2594
- Output Fields | 2594
- Sample Output | 2594
- Release Information | 2595

Syntax

```
clear dhcpv6 server statistics
<bulk-leasequery-connections>
<interface interface-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCPv6 local server bulk leasequery statistics.
logical-system *logical-system-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

Release Information

Command introduced in Junos OS Release 9.6.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show dhcpv6 server statistics | 2877

clear diameter function statistics

IN THIS SECTION

- Syntax | 2595
- Description | 2595
- Options | 2596
- Required Privilege Level | 2596
- Output Fields | 2596
- Sample Output | 2596
- Release Information | 2596

Syntax

```
clear diameter function <function-name> statistics
```

Description

Clear current statistics accumulated for a specified function (application) or for all functions associated with the Diameter instance.

Options

function-name (Optional) Clear statistics for the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions.

Required Privilege Level

clear

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear diameter function statistics

```
user@host> clear diameter function jsrc statistics
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

- Gx-Plus for Provisioning Subscribers Overview | 1393
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview | 1468
- show diameter | 2883
- show diameter function | 2892
- show diameter function statistics | 2898

clear diameter peer

IN THIS SECTION

- Syntax | 2597
- Description | 2597
- Options | 2597
- Required Privilege Level | 2597
- Output Fields | 2598
- Sample Output | 2598
- Release Information | 2598

Syntax

```
clear diameter peer peer-name
<connection | statistics>
```

Description

Delete the specified Diameter peer and clear all statistics or only current statistics for the specified peer.

Options

peer-name Delete the Diameter peer.

connection (Optional) Clear all peer statistics and restart the peer state machine for the specified Diameter peer. This is the default action.

statistics (Optional) Clear current statistics for the specified Diameter peer.

Required Privilege Level

clear

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear diameter peer

```
user@host> clear diameter peer peer5 connection
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- show diameter | 2883
- show diameter peer | 2914
- show diameter peer map | 2920
- show diameter peer statistics | 2924

clear extensible-subscriber-services counters

IN THIS SECTION

- Syntax | 2599
- Description | 2599
- Required Privilege Level | 2599
- Sample Output | 2599
- Release Information | 2599

Syntax

```
clear extensible-subscriber-services counters
```

Description

Clear values of all event counters to zero. Active sessions and services counters are not reset.

Required Privilege Level

view

Sample Output

command-name

```
#clear extensible-subscriber-services counters
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show extensible-subscriber-services counters

clear extensible-subscriber-services sessions

IN THIS SECTION

- Syntax | **2600**
- Description | **2600**
- Options | **2600**

- Required Privilege Level | 2600
- Sample Output | 2600
- Release Information | 2600

Syntax

```
clear extensible-subscriber-services sessions <accounting-session-id>
```

Description

Clear extensible subscriber service sessions in any state by executing the operational script or the application to remove all services created. Specify an accounting session ID to clear a specific session. If you do not specify an accounting session ID in the command, the command clears all sessions.

Options

accounting-session-id (Optional) Identifier of the ESSM session you want cleared.

Required Privilege Level

view

Sample Output

command-name

```
#clear extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show extensible-subscriber-services sessions

clear ipv6 router-advertisement

IN THIS SECTION

- Syntax | 2601
- Description | 2601
- Options | 2601
- Required Privilege Level | 2602
- Output Fields | 2602
- Sample Output | 2602
- Release Information | 2602

Syntax

```
clear ipv6 router-advertisement
<interface interface>
<logical-system (all | logical-system-name)>
```

Description

Clear IPv6 router advertisement counters.

Options

none	Clear IPv6 router advertisement counters for all interfaces.
interface *interface*	(Optional) Clear IPv6 router advertisement counters for the specified interface.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 router-advertisement

```
user@host> clear ipv6 router-advertisement
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show ipv6 router-advertisement | 2939

clear network-access aaa statistics

IN THIS SECTION

- Syntax | 2603
- Description | 2603
- Options | 2603
- Required Privilege Level | 2604

Output Fields | 2604
Sample Output | 2604
Release Information | 2607

Syntax

```
clear network-access aaa statistics
<accounting>
<address-assignment (client | pool pool-name)>
<authentication>
<dynamic-requests>
<radius>
<re-authentication>
<session-limit-per-username username username access-profile profile-name>
<terminate-code>
```

Description

Clear AAA statistics.

Options

accounting	(Optional) Clear AAA accounting statistics.
address-assignment client	(Optional) Clear AAA address-assignment statistics for the client.
address-assignment pool *pool-name*	(Optional) Clear AAA address-assignment pool statistics.
authentication	(Optional) Clear AAA authentication statistics.
dynamic-requests	(Optional) Clear AAA dynamic-request statistics.
radius	(Optional) Clears the values in the Peak and Exceeded columns only.
re-authentication	(Optional) Clear AAA reauthentication statistics.

session-limit-per-username (MX Series routers only) (Optional) Clear all blocked request statistics for all access profiles from the username session-limit table. You can also specify additional options:

- `username` *username*—Clear the blocked request statistics for the specified username across all access profiles. A given username can be used in more than one access profile.
- `access-profile` *profile-name*—Clear the blocked request statistics for all usernames in the specified access profile.

> **NOTE**: This command does not clear (delete) the entry in the session-limit table. Entries in the table are added or deleted during session login or logout processing.

terminate-code (Optional) Clear AAA termination code statistics.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access aaa statistics accounting

```
user@host> clear network-access aaa statistics accounting
```

clear network-access aaa statistics address-assignment pool

```
user@host> clear network-access aaa statistics address-assignment pool isp_1
```

clear network-access aaa statistics radius

```
user@host> clear network-access aaa statistics radius
```

clear network-access aaa statistics session-limit-per-username (All Usernames Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
pqr@example.net    BNG2               3                    4

user@host> clear network-access aaa statistics session-limit-per-username
user@host> show network-access aaa statistics on-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               0                    5
abc@example.net    BNG2               0                    5
pqr@example.net    BNG2               0                    4
```

clear network-access aaa statistics session-limit-per-username (Specific Username Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
rkv@example.net    BNG2               2                    5
pqr@example.net    BNG2               3                    4
user@host> clear network-access aaa statistics session-limit-per-username username
rkv@example.net
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               3                    5
```

```
rkv@example.net    BNG2              0                   5
pqr@example.net    BNG2              3                   4
```

clear network-access aaa statistics session-limit-per-username (All Usernames for Specific Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                   4
xyz@example.net    BNG1               3                   5
rkv@example.net    BNG2               2                   5
pqr@example.net    BNG2               3                   4
user@host> clear network-access aaa statistics session-limit-per-username access-profile BNG2
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                   4
xyz@example.net    BNG1               3                   5
123@example.net    BNG2               0                   5
pqr@example.net    BNG2               0                   4
```

clear network-access aaa statistics session-limit-per-username (Specific Username in Specific Access Profile)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                   4
xyz@example.net    BNG1               3                   5
rkv@example.net    BNG2               2                   5
pqr@example.net    BNG2               3                   4
user@host> clear network-access aaa statistics session-limit-per-username username
rkv@example.net access-profile BNG2
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               3                   4
xyz@example.net    BNG1               3                   5
```

```
rkv@example.net     BNG2            0                   5
pqr@example.net     BNG2            3                   4
```

Release Information

Command introduced in Junos OS Release 10.0.

`radius` option introduced in Junos OS Release 11.4

`terminate-code` option introduced in Junos OS Release 11.4.

`session-limit-per-username` option introduced in Junos OS Release 18.4R1 on MX Series routers.

RELATED DOCUMENTATION

- Verifying and Managing Subscriber AAA Information
- Understanding Session Options for Subscriber Access
- Limiting the Number of Active Sessions per Username and Access Profile
- show network-access aaa statistics

clear network-access aaa subscriber

IN THIS SECTION

- Syntax | **2608**
- Description | **2608**
- Options | **2608**
- Required Privilege Level | **2608**
- Output Fields | **2609**
- Sample Output | **2609**
- Release Information | **2609**

Syntax

```
clear network-access aaa subscriber
<session-id identifier <reconnect>>
<statistics username username>
<username username <reconnect>>
```

Description

Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.

Options

reconnect (Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16). You can apply this option to either a Layer 2 wholesale session or a conventionally auto-sensed dynamic VLAN supporting a PPPoE session.

In the latter case, this option triggers a PPPoE session logout and removal of the dynamic VLAN logical interface. This is followed by authorization of the access-line to attempt creation of a dynamic VLAN IFL supporting Layer 2 wholesale session in its place.

session-id *identifier* (Optional) Log out the subscriber based on the subscriber session identifier.

statistics username *username* (Optional) Clear AAA subscriber statistics and log out the subscriber.

username *username* (Optional) Log out the AAA subscriber.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber statistics username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

clear network-access aaa subscriber session-id (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber session-id 1
```

Release Information

Command introduced in Junos OS Release 9.1.

`reconnect` and `session-id` options added in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

clear network-access gx-plus replay

IN THIS SECTION

- Syntax | 2610
- Description | 2610
- Options | 2610
- Required Privilege Level | 2611
- Output Fields | 2611
- Sample Output | 2611
- Release Information | 2611

Syntax

```
clear network-access gx-plus replay
```

Description

Clear pending Gx-Plus login and logout requests (replays). Sends 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 returns a JDER message to initiate discovery of all subscribers. When this discovery completes, all pending subscriber requests are cleared.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access gx-plus replay

```
user@host> clear network-access gx-plus replay
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear network-access gx-plus statistics | 2611

show network-access gx-plus | 3014

clear network-access gx-plus statistics

IN THIS SECTION

- Syntax | 2612
- Description | 2612
- Options | 2612
- Required Privilege Level | 2612
- Output Fields | 2612
- Sample Output | 2612

- Release Information | 2612

Syntax

```
clear network-access gx-plus statistics
```

Description

Clear Gx-Plus statistics.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access gx-plus statistics

```
user@host> clear network-access gx-plus statistics
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

show network-access gx-plus | 3014

clear network-access ocs statistics

IN THIS SECTION

- Syntax | 2613
- Description | 2613
- Required Privilege Level | 2613
- Output Fields | 2613
- Sample Output | 2614
- Release Information | 2614

Syntax

```
clear network-access ocs statistics
```

Description

Clear Online Charging System (OCS) provisioning statistics information.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear network-access ocs statistics

```
user@host> clear network-access ocs statistics
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- ocs (Diameter Applications) | 2090
- show network-access ocs | 3037
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | 1413
- Understanding Gy Interactions Between the Router and the OCS | 1431
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439

clear network-access pcrf

IN THIS SECTION

- Syntax | 2615
- Description | 2615
- Options | 2615
- Required Privilege Level | 2615
- Output Fields | 2615
- Sample Output | 2615
- Sample Output | 2615
- Release Information | 2616

Syntax

```
clear network-access pcrf (statistics | subscribers)
```

Description

Clear Policy and Charging Rules Function (PCRF) provisioning statistics and subscribers information.

Options

statistics (Optional) Clear PCRF provisioning statistics.

subscribers (Optional) Force logout of all PCRF subscribers. Used with draining mode to flush all subscribers before making substantial configuration changes.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear network-access pcrf statistics

```
user@host> clear network-access pcrf statistics
```

Sample Output

clear network-access pcrf subscribers

```
user@host> clear network-access pcrf subscribers
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- pcrf (Diameter Applications) | 2174
- show network-access pcrf | 3040
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

clear system subscriber-management statistics

IN THIS SECTION

- Syntax | 2616
- Description | 2616
- Options | 2617
- Required Privilege Level | 2617
- Output Fields | 2617
- Sample Output | 2617
- Release Information | 2617

Syntax

```
clear system subscriber-management statistics
```

Description

Clear subscriber-management statistics.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear subscriber-management statistics

```
user@host> clear subscriber-management statistics
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- show system subscriber-management statistics | 3134
- clear system subscriber-management arp
- clear system subscriber-management ipv6 neighbors

request ancp oam interface

IN THIS SECTION

- Syntax | 2618
- Description | 2618

- Options | 2618
- Required Privilege Level | 2618
- Output Fields | 2619
- Sample Output | 2619
- Release Information | 2619

Syntax

```
request ancp oam interface
(interface-name | interface-set set-name)
<count count>
<timeout duration>
```

Description

Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify either an ANCP interface or an ANCP interface set. The access node responds to the NAS with the results of the test.

Options

interface-name	Name of the ANCP interface on whose local loop the loopback test is run.
interface-set *set-name*	Name of the ANCP interface set on whose local loop the loopback test is run.
count *count*	(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.
timeout *duration*	(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam interface

```
user@host> request ancp oam interface ge-1/0/4.12 count 5 timeout 40
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Test the Local Loop | 1326

request ancp oam neighbor

IN THIS SECTION

- Syntax | 2620
- Description | 2620
- Options | 2620
- Required Privilege Level | 2620
- Output Fields | 2621
- Sample Output | 2621

Release Information | 2621

Syntax

```
request ancp oam neighbor
(ip-address ip-address | system-name neighbor-name)
subscriber identifier-string
<count count>
<timeout duration>
```

Description

Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify both the access node and the subscriber. The access node responds to the NAS with the results of the test.

Options

ip-address *ip-address*	IP address that specifies the access node on whose local loop the loopback test is run.
system-name *neighbor-name*	System name that specifies the access node on whose local loop the loopback test is run.
subscriber *identifier-string*	Access identifier that specifies the subscriber on whose local loop the loopback test is run.
count *count*	(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.
timeout *duration*	(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam subscriber

```
user@host> request ancp oam neighbor 203.0.113.21 subscriber “dslam port-1-11”
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Test the Local Loop | 1326

request ancp oam port-down

IN THIS SECTION

- Syntax | 2622
- Description | 2622
- Options | 2622
- Required Privilege Level | 2623
- Output Fields | 2623
- Sample Output | 2623

- Release Information | 2623

Syntax

```
request ancp oam port-down
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Description

Simulate an ANCP Port Down message on the specified access loop for troubleshooting or to mitigate an abnormal condition. Triggers removal of the corresponding out-of-band triggered, autosensed dynamic VLAN session for which no ANCP-sourced information exists. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port-Up message, meaning that you cannot use this command to initiate a Port Down condition when the access node has already reported a Port Up condition.

Options

aci	ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.
ari	ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.
ip-address	IP address that specifies the access node from which the message is simulated.
physical-interface-name	Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.
vlan-id	ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor detail`, `show subscribers client-type vlan-oob detail`, and the `show subscribers summary` commands before and after initiating the Port Down message to verify the operation.

Sample Output

request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-down neighbor 192.168.25.31 circuit-id line-aci-1 remote-id
line-ari-1 outer-vlan-id 126
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

request ancp oam port-up

IN THIS SECTION

- Syntax | 2624
- Description | 2624
- Options | 2624
- Required Privilege Level | 2625

- Output Fields | **2625**
- Sample Output | **2625**
- Release Information | **2625**

Syntax

```
request ancp oam port-up
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Description

Simulate an ANCP Port Up message on the specified access loop for troubleshooting or to mitigate an abnormal condition. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port Down message, meaning that you cannot use this command to initiate a Port Up condition when the access node has already reported a Port Down condition.

Options

aci	ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.
ip-address	IP address that specifies the access node from which the message is simulated.
ari	ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.
physical-interface-name	Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.
vlan-id	ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor detail`, `show subscribers client-type vlan-oob detail`, and the `show subscribers summary` commands before and after initiating the Port Up message to verify the operation.

Sample Output

request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-up neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-
ari-1 outer-vlan-id 126
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

request dhcp relay bulk-leasequery

IN THIS SECTION

- Syntax | 2626
- Description | 2626
- Options | 2626
- Required Privilege Level | 2627

Output Fields | 2627
Sample Output | 2627
Release Information | 2628

Syntax

```
request dhcp relay bulk-leasequery
<all-configured-ip>
<client-id | ipv4-address | mac-address>
<relay-id relay-id>
<remote-id remote-id>
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a bulk leasequery operation to update the binding information for multiple subscribers by requesting that the DHCPv4 relay agent send the bulk leasequery message to the configured DHCPv4 servers. The DHCP bulk leasequery feature must be configured on the DHCPv4 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv4 servers known to the relay agent.

Options

NOTE: 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*.

all-configured-ip Update 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 and regardless of remote Id or relay ID. This enables the relay agent to update its database with all address changes that occurred after some point in time.

client-id	Update binding information for clients identified by the specified DHCPv4 client option (option 61).
ipv4-address	Update binding information for the most recent client that was assigned that IP address.
logical-system *logical-system-name*	Specify an optional logical system for the DHCP server being queried. The default logical system is used by default.
mac-address	Update binding information for the most recent client that has that MAC address.
relay-id *relay-id*	Update binding information for all currently active leases assigned to the client that has the specified Relay Agent Identifier suboption (suboption 12) of the DHCP relay agent information option (option 82). Suboption 12 uniquely identifies a relay agent within an administrative domain, which consists of all DHCP servers and relay agents that communicate with each other.
remote-id *remote-id*	Update binding information for all currently active leases assigned to clients that use that Agent Remote ID (suboption 2) of the DHCP relay agent information option (option 82).
routing-instance *routing-instance-name*	Specify an optional routing instance for the DHCP server being queried. The default routing instance is used by default.
server-address *server-address*	Specify the address of a DHCP local server to query.
server-group *group-name*	Specify the name of a group of DHCP local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCP relay agent initiates the bulk leasequery operation.

Sample Output

request dhcp relay bulk-leasequery

```
user@host> request dhcp relay bulk-leasequery 4/1224 server-address 192.168.10.10
```

Release Information

Command introduced in Junos OS Release 16.1.

`all-configured-ip`, `relay-id`, and `remote-id` options added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | 824

request dhcp relay leasequery

IN THIS SECTION

- Syntax | 2628
- Description | 2629
- Options | 2629
- Required Privilege Level | 2629
- Output Fields | 2630
- Sample Output | 2630
- Release Information | 2630

Syntax

```
request dhcp relay leasequery (client-id | ipv4-address | mac-address)
gateway-address giaddr;
<server-address address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a leasequery operation to update the binding information for a specific subscriber by requesting that DHCPv4 relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured on the DHCPv4 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv4 servers known to the relay agent.

Options

NOTE: You must configure at least one of *client-id*, *ipv4-address*, or *mac-address*.

client-id	Use the DHCPv4 client identifier option (option 61) to identify the client whose binding information is requested.
gateway-address *giaddr*	Gateway address of the relay agent making the request; the value in the giaddr field of the DHCP packet. You must always specify a gateway address in your request.
ipv4-address	Use the IPv4 address to identify the client whose binding information is requested.
logical-system *logical-system-name*	Specify an optional logical system for the DHCP servers being queried. The default logical system is used by default.
mac-address	Use the MAC address to identify the client whose binding information is requested.
routing-instance *routing-instance-name*	Specify an optional routing instance for the DHCPv4 servers being queried. The default routing instance is used by default.
server-address *server-address*	Specify the IP address of the DHCPv4 local server to query.
server-group *group-name*	Specify the name of a group of DHCPv4 local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCP relay agent initiates the leasequery operation.

Sample Output

request dhcp relay leasequery

```
user@host> request dhcp relay leasequery 192.168.25.25 server-group DHCPgroup-10
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | 824

request dhcp server reconfigure

IN THIS SECTION

- Syntax | 2631
- Description | 2631
- Options | 2631
- Required Privilege Level | 2632
- Output Fields | 2632
- Sample Output | 2632
- Release Information | 2632

Syntax

```
request dhcp server reconfigure (all | address | interface interface-name | logical-system logical-system-name | routing-instance routing-instance-name)
```

Description

Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcp server binding` command.

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 (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all — Initiate reconfiguration for all DHCP clients.

address — Initiate reconfiguration for DHCP client with the specified IP address or MAC address.

interface *interface-name* — Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).

> **NOTE**: You cannot use the `interface interface-name` option with the `request dhcp server reconfigure` command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.

logical-system *logical-system-name* — Initiate reconfiguration for all DHCP clients on the specified logical system.

routing-instance *routing-instance-name* — Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request dhcp server reconfigure

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

request dhcpv6 server reconfigure

IN THIS SECTION

- Syntax | 2633
- Description | 2633
- Options | 2633
- Required Privilege Level | 2634

- Output Fields | 2634
- Sample Output | 2634
- Release Information | 2634

Syntax

```
request dhcpv6 server reconfigure (all | address | client-id | interface interface-name | logical-system
logical-system-name | routing-instance routing-instance-name | session-id)
```

Description

Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcpv6 server binding` command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all	Initiate reconfiguration for all DHCPv6 clients.
address	Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.
client-id	Initiate reconfiguration for DHCPv6 client with the specified client ID.
interface ***interface-name***	Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).
logical-system ***logical-system-name***	Initiate reconfiguration for all DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name*	Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.
session-id	Initiate reconfiguration for DHCPv6 client with the specified session ID.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001db8::2/16
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview | 874

request dhcpv6 relay bulk-leasequery

IN THIS SECTION

- Syntax | 2635
- Description | 2635

- Options | 2635
- Required Privilege Level | 2636
- Output Fields | 2636
- Sample Output | 2636
- Release Information | 2636

Syntax

```
request dhcpv6 relay bulk-leasequery
<client-id | ipv6-prefix| link-address ipv6-link-address | relay-id relay-id | remote-id remote-id>
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a bulk leasequery operation to update the binding information for multiple subscribers by requesting that the DHCPv6 relay agent send the bulk leasequery message to the configured DHCPv6 servers. The DHCP bulk leasequery feature must be configured on the DHCPv6 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv6 servers known to the relay agent.

Options

NOTE: 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*.

client-id	Update binding information for clients identified by the DHCP Unique Identifier (DUID) specified DHCPv6 Client ID option (option 1).
ipv6-prefix	Update binding information for clients identified by the specified IPv6 prefix.

link-addr *ipv6-link-address*	Update binding information for clients associated with the specified IPv6 network segment. The link address is an address that a relay agent may have used in a Relay-Forward message.
logical-system *logical-system-name*	Specify an optional logical system for the clients. The default logical system is used by default.
relay-id *relay-id*	Update binding information for clients associated with the DHCPv6 relay identified by the DHCP Unique Identifier (DUID) in the specified Relay-ID option (option 53).
remote-id *remote-id*	Update binding information for clients associated with the specified Relay Agent Remote-ID option (option 37).
routing-instance *routing-instance-name*	Specify an optional routing instance for the clients. The default routing instance is used by default.
server-address *server-address*	Specify the IPv6 address of the DHCP local server to query.
server-group *group-name*	Specify the name of a group of DHCP local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCPv6 relay agent initiates the bulk leasequery operation.

Sample Output

request dhcpv6 relay bulk-leasequery

```
user@host> request dhcpv6 relay bulk-leasequery LL0x1-00:00:65:03:01:02 server-address
2001:db8:3000:0:8001::5/128
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | 824

request dhcpv6 relay leasequery

IN THIS SECTION

- Syntax | 2637
- Description | 2637
- Options | 2638
- Required Privilege Level | 2638
- Output Fields | 2638
- Sample Output | 2638
- Release Information | 2638

Syntax

```
request dhcpv6 relay leasequery (client-id | ipv6-prefix)
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a leasequery operation to update the binding information for a specific subscriber by requesting that the DHCPv6 relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured on the DHCPv6 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCP servers known to the relay agent.

Options

NOTE: You must configure at least one of *client-id* or *ipv6-prefix*.

client-id	Use the DHCPv6 Client ID option (option 1) to identify the client whose binding information is requested.
ipv6-prefix	Use the IPv6 prefix to identify the client whose binding information is requested.
logical-system *logical-system-name*	Specify an optional logical system for the DHCPv6 servers being queried. The default logical system is used by default.
routing-instance *routing-instance-name*	Specify an optional routing instance for the DHCPv6 servers being queried. The default routing instance is used by default.
server-address *server-address*	Specify the IPv6 address of the DHCPv6 local server to query.
server-group *group-name*	Specify the name of a group of DHCPv6 local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCPv6 relay agent initiates the leasequery operation.

Sample Output

request dhcpv6 relay leasequery

```
user@host> request dhcpv6 relay leasequery 2001:db8:3000:0:8001::5/128
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database | 824

request network-access aaa accounting

IN THIS SECTION

- Syntax | 2639
- Description | 2639
- Options | 2639
- Required Privilege Level | 2640
- Sample Output | 2640
- Release Information | 2640

Syntax

```
request network-access aaa accounting (baseline | suspend | resume)
```

Description

Suspend accounting processes; determine a baseline of the statistical details while accounting is suspended; and restart accounting operations after baselining is completed. This command is useful in service provider environments when an upgrade of the server infrastructure is critical and needed immediately. 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. While accounting is suspended, subscribers can continue to log in and log out.

Options

baseline (Optional) Determine a baseline of accounting statistics for current subscriber sessions. Applies to only those subscribers for which interim accounting is configured. The router implements the baseline by reading and storing the statistics at the time the baseline is set

and then subtracting this baseline when you retrieve baseline-relative statistics after accounting resumes.

resume Restart the accounting processes for all logged-in subscriber sessions after baselining of statistics completes.

suspend Temporarily halt accounting processes for all logged-in subscriber sessions.

Required Privilege Level

view

Sample Output

request network-access aaa accounting suspend

```
user@host> request network-access aaa accounting suspend
```

request network-access aaa accounting baseline

```
user@host> request network-access aaa accounting baseline
```

request network-access aaa accounting resume

```
user@host> request network-access aaa accounting resume
```

Release Information

Command introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

Configuring RADIUS Accounting Suspension and Baselining Accounting Statistics | 599

Suspending RADIUS Accounting and Baselining Accounting Statistics Overview | 595

request network-access aaa replay pending-accounting-stops

IN THIS SECTION

- Syntax | 2641
- Description | 2641
- Options | 2641
- Required Privilege Level | 2641
- Output Fields | 2641
- Sample Output | 2642
- Release Information | 2642

Syntax

```
request network-access aaa replay pending-accounting-stops
```

Description

Force the router to attempt contact with the accounting sever immediately, rather than allowing it to wait until the periodic interval has expired.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request network-access aaa replay pending-accounting-stops

```
user@host> request network-access aaa replay pending-accounting-stops
replay started
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Forcing the Router to Contact the Accounting Server Immediately | 592

show accounting pending-accounting-stops | 2669

request network-access aaa subscriber modify session-id

IN THIS SECTION

- Syntax | 2643
- Description | 2643
- Options | 2643
- Required Privilege Level | 2643
- Output Fields | 2643
- Sample Output | 2644
- Release Information | 2644

Syntax

```
request network-access aaa subscriber modify session-id subscriber-session-id predefined-variable variable-option
```

Description

Modify a predefined variable that is applied to a subscriber who is currently logged in to the network.

Options

predefined-variable	Name of the predefined variable that you want to modify.
subscriber-session-id	ID of the subscriber session.
variable-option	Name of the variable option that you want to apply to the predefined variable.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 94 on page 2643 lists possible messages that might be returned.

Table 94: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Successful completion`	Variable was successfully modified	–
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Sample Output

request network-access aaa subscriber modify session-id

```
user@host> request network-access aaa subscriber modify session-id 49 junos-cos-traffic-control-profile TCP-gold
Successful completion
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

CLI-Activated Subscriber Services

request network-access aaa subscriber set session-id

IN THIS SECTION

- Syntax | 2645
- Description | 2645
- Options | 2645
- Required Privilege Level | 2645
- Output Fields | 2645
- Sample Output | 2646
- Release Information | 2646

Syntax

```
request network-access aaa subscriber set session-id subscriber-session-id provisioning-state
none
```

Description

Release control of the PCRF over the specified subscriber session. In response, AAA clears the subscriber's provisioning state and sends a terminated request to the PCRF indicating the subscriber is no longer available.

Options

subscriber-session-id	ID of the subscriber session.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 95 on page 2645 lists possible error messages that might be returned if the service activation fails.

Table 95: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
`Service activation/ deactivation already in progress`	Another service activation/ deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.

Table 95: Service Activation/Deactivation Error Messages *(Continued)*

Message	Description	Corrective Action
`Session identifier is not for a subscriber session`	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber set session-id

```
user@host> request network-access aaa subscriber set session-id session-id 49 provisioning-state
none
Successful completion
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Disabling PCRF Control of a Subscriber Session | 1407

Local and Remote Service Activation and Deactivation Using the CLI

request services extensible-subscriber-services reload-dictionary

IN THIS SECTION

- Syntax | 2647
- Description | 2647
- Options | 2647
- Required Privilege Level | 2647

Sample Output | 2647
Release Information | 2647

Syntax

```
request services extensible-subscriber-services reload-dictionary
```

Description

Reload the configured dictionary to essmd.

Options

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Sample Output

command-name

```
# request services extensible-subscriber-services reload-dictionary
Dictionary reloaded successfully
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dictionary | 1791

show extensible-subscriber-services dictionary

show extensible-subscriber-services dictionary attributes

show extensible-subscriber-services dictionary services

Understanding the Dictionary File

request services static-subscribers login group

IN THIS SECTION

- Syntax | 2648
- Description | 2648
- Options | 2648
- Required Privilege Level | 2649
- Output Fields | 2649
- Sample Output | 2649
- Release Information | 2649

Syntax

```
request services static-subscribers login group group-name
```

Description

Resets the state of an interface group on which static subscribers were forcibly logged out by the `request services static-subscribers logout group` command. This action enables static subscriber to login on the interfaces in the group.

Options

group *group-name* Group of static subscriber interfaces on which static subscribers have been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers login group

```
user@host> request services static-subscribers login group boston
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Resetting the State of an Interface Group for Static Subscriber Login | 1513

request services static-subscribers logout group | 2649

request services static-subscribers logout group

IN THIS SECTION

- Syntax | 2650
- Description | 2650
- Options | 2650
- Required Privilege Level | 2650
- Output Fields | 2650
- Sample Output | 2650

Release Information | 2650

Syntax

```
request services static-subscribers logout group igroup-name
```

Description

Force static subscribers on the interfaces in the group to be logged out. No subscriber can subsequently log in on the interface group until the interface state is reset by a router reset or the `request services static-subscribers login group` command.

Options

group *group-name* Group of static subscriber interfaces on which static subscribers have been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers logout group

```
user@host> request services static-subscribers logout group boston
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Forcing a Group of Static Subscribers to Be Logged Out | 1513

request services static-subscribers login group | 2648

request services static-subscribers login interface

IN THIS SECTION

- Syntax | 2651
- Description | 2651
- Options | 2651
- Required Privilege Level | 2652
- Output Fields | 2652
- Sample Output | 2652
- Release Information | 2652

Syntax

```
request services static-subscribers login interface interface-name
```

Description

Resets the state of an interface on which a static subscriber was forcibly logged out by the `request services static-subscribers logout interface` command. This action enables a static subscriber to login on the interface.

Options

interface *interface-name* Static interface on which a static subscriber has been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers login interface

```
user@host> request services static-subscribers login interface ge-2/0/1.5
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Resetting the State of an Interface for Static Subscriber Login | 1512

request services static-subscribers logout interface | 2652

request services static-subscribers logout interface

IN THIS SECTION

- Syntax | 2653
- Description | 2653
- Options | 2653
- Required Privilege Level | 2653
- Output Fields | 2653
- Sample Output | 2653

Release Information | 2653

Syntax

```
request services static-subscribers logout interface interface-name
```

Description

Force static subscriber on the interface to be logged out. No subscriber can subsequently log in on the interface until the interface state is reset by a router reset or the `request services static-subscribers login interface` command.

Options

interface *interface-name* Static interface on which a static subscriber has been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers logout interface

```
user@host> request services static-subscribers logout interface ge-2/0/1.5
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Forcing a Static Subscriber to Be Logged Out | 1512

request services static-subscribers login interface | 2651

request services subscribers

IN THIS SECTION

- Syntax | 2654
- Description | 2654
- Options | 2654
- Required Privilege Level | 2655
- Sample Output | 2655
- Release Information | 2655

Syntax

```
request services subscribers set subscriber-profile profile client-id client-id
```

Description

Set the subscriber profile associated with the given subscriber.

Options

profile Name of the subscriber profile to create or override the currently active subscriber profile for the given subscriber.

client-id Client session ID assigned to the subscriber.

Required Privilege Level

view

Sample Output

request services subscriber set subscriber-profile tc_act_prof client-id

```
user@host> request services subscriber set subscriber-profile tc_act_prof client-id
2533274790395909 | display xml
rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1I0/junos"
    packet-triggered-subscribers-information xmlns="http://xml.juniper.net/junos/11.1I0/junos-
packet-triggered-subscribers"
        service-subscribers-request-result junos:style="success"
        /service-subscribers-request-result
    /packet-triggered-subscribers-information
    cli
        banner/banner
    /cli
/rpc-reply
```

Release Information

Command introduced in Junos OS Release 11.4.

request services subscribers clear

IN THIS SECTION

- Syntax | **2656**
- Description | **2656**
- Options | **2656**
- Required Privilege Level | **2656**
- Sample Output | **2656**
- Release Information | **2656**

Syntax

```
request services subscribers clear subscriber-profile profile client-id client-id
```

Description

Clear the subscriber profile associated with the given subscriber.

Options

profile Name of the subscriber profile to clear the active subscriber profile for the given subscriber.

client-id Client session ID assigned to the subscriber.

Required Privilege Level

clear

Sample Output

request services subscriber clear subscriber-profile tc_act_prof client-id

```
user@host>request services subscriber clear subscriber-profile tc_act_prof client-id
2533274790395909 | display xml
rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1I0/junos"
    packet-triggered-subscribers-information xmlns="http://xml.juniper.net/junos/11.1I0/junos-
packet-triggered-subscribers"
        service-subscribers-request-result junos:style="success"
        /service-subscribers-request-result
    /packet-triggered-subscribers-information
    cli
        banner/banner
    /cli
/rpc-reply
```

Release Information

Command introduced in Junos OS Release 11.4.

request system reboot (Junos OS)

IN THIS SECTION

- Syntax | 2657
- Syntax (EX Series Switches and EX Series Virtual Chassis) | 2658
- Syntax (MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis) | 2658
- Syntax (QFabric Systems) | 2658
- Syntax (QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric) | 2659
- Syntax (TX Matrix Router) | 2659
- Syntax (TX Matrix Plus Router) | 2659
- Description | 2659
- Options | 2660
- Additional Information | 2663
- Required Privilege Level | 2664
- Output Fields | 2664
- Sample Output | 2664
- Release Information | 2667

Syntax

```
request system reboot
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
<other-routing-engine>
```

Syntax (EX Series Switches and EX Series Virtual Chassis)

```
request system reboot
<all-members | local | member member-id>
<at time>
<in minutes>
<media (external | internal)> | <media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
<slice slice>
```

Syntax (MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis)

```
request system reboot
<all-members | local | member member-id>
<at time>
<both-routing-engines>
<in minutes>
<media (external | internal)> | <media (compact-flash | disk | usb)> | <junos | network | oam |
usb>
<message "text">
<other-routing-engine>
```

Syntax (QFabric Systems)

```
request system reboot
<all <graceful>>
<at time>
<director-device name>
<director-group <graceful>>
<fabric <graceful>>
<in minutes>
<in-service>
<media>
<message “text”>
<node-group name>
<slice slice>
```

Syntax (QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric)

```
request system reboot
<all-members | local | member member-id>
<at time>
<in minutes>
<in-service>
<hypervisor>
<junos | network | oam | usb>
<message “text”>
<slice slice>
```

Syntax (TX Matrix Router)

```
request system reboot
<all-chassis | all-lcc | lcc number /  scc>
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
```

Syntax (TX Matrix Plus Router)

```
request system reboot
<all-chassis | all-lcc | lcc number | sfc number>
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
<partition (1 | 2 | alternate)>
```

Description

Use this command to reboot the device software.

This command can be used on standalone devices and on devices supported in a Virtual Chassis, Virtual Chassis Fabric, or QFabric system.

Starting with Junos OS Release 15.1F3, the `request system reboot` command reboots only the guest operating system on the PTX5000 with RE-PTX-X8-64G and, MX240, MX480, and MX960 with RE-S-X6-64G.

Starting with Junos OS Release 15.1F5, the `request system reboot` command reboots only the guest operating system on the MX2010, and MX2020 with REMX2K-X8-64G.

Starting from Junos OS Release 17.2R1, PTX10008 routers do not support the `request system reboot` command. Starting from Junos OS Release 17.4R1, PTX10016 routers do not support the `request system reboot` command. Use the `request vmhost reboot` command instead of the `request system reboot` command on the PTX10008 and PTX10016 routers to reboot the Junos OS software package or bundle on the router. See request vmhost reboot.

Starting from Junos OS Release 19.1R1, the PTX10002-60C router and the QFX10002-60C switch do not support the `request system reboot` command. Use the `request vmhost reboot` command instead of the `request system reboot` command on these devices to reboot the Junos OS software package or bundle on the device. See request vmhost reboot.

On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Options

The options described here are not all supported on every platform or release of Junos OS. Refer to the Syntax sections for the options commonly available on each type of platform.

none — Reboot the software immediately.

all-chassis — (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.

all-lcc — (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.

all-members | local | member ***member-id*** — (Optional) Specify which member of the Virtual Chassis to reboot:

- `all-members`—Reboots each switch that is a member of the Virtual Chassis.
- `local`—Reboots only the local switch (switch where you are logged in).
- `member` *member-id*—Reboots the specified member switch of the Virtual Chassis

at *time* (Optional) Time at which to reboot the software, specified in one of the following ways:

- `now`—Stop or reboot the software immediately. This is the default.
- *`+minutes`*—Number of minutes from now to reboot the software.
- *`yymmddhhmm`*—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- *`hh:mm`*—Absolute time on the current day at which to stop the software, specified in 24-hour time.

both-routing-engines (Optional) Reboot both Routing Engines at the same time.

fast-boot (Optional, QFX Series) Minimizes traffic loss and downtime of network ports by leaving the network ports up during the system reboot.

hypervisor (Optional) Reboot Junos OS, host OS, and any installed guest VMs.

in *minutes* (Optional) Number of minutes from now to reboot the software. The minimum value is 1. This option is an alias for the `at` *`+minutes`* option.

in-service (Optional) Enables you to reset the software state (no software version change) of the system with minimal disruption in data and control traffic.

junos (Optional) Reboot from the Junos OS (main) volume.

lcc *number* —(Optional) Line-card chassis (LLC) number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

media (compact-flash | disk | (Optional) Use the indicated boot medium for the next boot.

removable-compact-flash | usb)

media (external | internal) (Optional) Use the indicated boot medium for the next boot:

- `external`—Reboot the device using a software package stored on an external boot source, such as a USB flash drive.
- `internal`—Reboot the device using a software package stored in an internal memory source.

message "*text*" (Optional) Message to display to all system users before stopping or rebooting the software.

network (Optional) Reboot using the Preboot Execution Environment (PXE) boot method over the network.

oam (Optional) Reboot from the maintenance volume (OAM volume, usually the compact flash drive).

other-routing-engine (Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is rebooted.

partition *partition* (Optional) Reboot using the specified partition on the boot media. This option is equivalent to the `slice` option that is supported on some devices. Specify one of the following *partition* values:

- `1`—Reboot from partition 1.
- `2`—Reboot from partition 2.
- `alternate`—Reboot from the alternate partition.

scc (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

sfc *number* (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace *number* with 0.

slice *slice* (Optional) Reboot using the specified partition on the boot media. This option was originally the `partitiion` option but was renamed to `slice` on EX Series and QFX Series switches. Specify one of the following *slice* values:

- 1—Reboot from partition 1.
- 2—Reboot from partition 2.
- alternate—Reboot from the alternate partition (which did not boot the switch at the last bootup).

> **NOTE**: The slice option is not supported on QFX Series switches that have no alternate slice when Junos OS boots as a Virtual Machine (VM). To switch to the previous version of Junos OS, issue the request system software rollback command.

usb (Optional) Reboot from a USB device.

The following options are available only on QFabric Systems:

all (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.

director-device *name* (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

node-group *name* (Optional) Reboots the software on a server Node group or a network Node group.

graceful (Optional) Enables the QFabric component to reboot with minimal impact to network traffic. This sub-option is only available for the all, fabric, anddirector-group options.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the show log command (see *show log*). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the show system processes command (see *show system processes*).

On a TX Matrix or TX Matrix Plus router, if you issue the request system reboot command on the primary Routing Engine, all the primary Routing Engines connected to the routing matrix are rebooted. If you

issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.

NOTE: Before issuing the `request system reboot` command on a TX Matrix Plus router with no options or the `all-chassis`, `all-lcc`, `lcc` *number*, or `sfc` options, verify that primary Routing Engine for all routers in the routing matrix are in the same slot number. If the primary Routing Engine for a line-card chassis is in a different slot number than the primary Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the `request system reboot` command.

NOTE: To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the primary Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
```

```
*** System shutdown message from root@test.example.net ***
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
 user@host> request system reboot in 120
 user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot in-service

```
user@switch> request system reboot in-service
Reboot the system ? [yes,no] yes
[Feb 22 02:37:04]:ISSU: Validating Image

 PRE ISSR CHECK:
 ---------------
 PFE Status                          : Online
 Member Id zero                      : Valid
```

```
 VC not in mixed or fabric mode        : Valid
 Member is single node vc              : Valid
 BFD minimum-interval check done       : Valid
 GRES enabled                          : Valid
 NSR enabled                           : Valid
 drop-all-tcp not configured           : Valid
 Ready for ISSR                        : Valid

warning: Do NOT use /user during ISSR. Changes to /user during ISSR may get lost!
Current image is jinstall-jcp-i386-flex-18.1.img
[Feb 22 02:37:14]:ISSU: Preparing Backup RE
Prepare for ISSR
[Feb 22 02:37:19]:ISSU: Backup RE Prepare Done
Spawning the backup RE
Spawn backup RE, index 1 successful
Starting secondary dataplane
Second dataplane container started
GRES in progress
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade for ISSR
Chassis ISSU Started
[Feb 22 02:42:55]:ISSU: Preparing Daemons
[Feb 22 02:43:00]:ISSU: Daemons Ready for ISSU
[Feb 22 02:43:05]:ISSU: Starting Upgrade for FRUs
[Feb 22 02:43:15]:ISSU: FPC Warm Booting
[Feb 22 02:44:16]:ISSU: FPC Warm Booted
[Feb 22 02:44:27]:ISSU: Preparing for Switchover
[Feb 22 02:44:31]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
Send ISSR done to chassisd on backup RE
Chassis ISSU Completed
Removing dcpfe0 eth1 128.168.0.16 IP
Bringing down bme00
Post Chassis ISSU processing done
[Feb 22 02:44:33]:ISSU: IDLE
Stopping primary dataplane
Clearing ISSU states
Console and management sessions will be disconnected. Please login again.
```

```
device_handoff successful ret: 0
Shutdown NOW!
[pid 14305]

*** FINAL System shutdown message from root@sw-duckhorn-01 ***

System going down IMMEDIATELY
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `other-routing-engine` introduced in Junos OS Release 8.0.

Option `sfc` introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Option `partition` changed to `slice` in Junos OS Release 10.0 for EX Series switches.

Option `both-routing-engines` introduced in Junos OS Release 12.1.

Option `fast-boot` introduced in Junos OS Release 14.1X53-D10 for QFX Series.

RELATED DOCUMENTATION

request system halt

Routing Matrix with a TX Matrix Plus Router Solutions Page

request vmhost reboot

clear system reboot

restart extensible-subscriber-services

IN THIS SECTION

- Syntax | **2668**
- Description | **2668**
- Required Privilege Level | **2668**
- Sample Output | **2668**

- Release Information | 2668

Syntax

```
restart extensible-subscriber-services
```

Description

Restart essmd.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Sample Output

command-name

```
# restart extensible-subscriber-services
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

disable (Extensible Subscriber Services Manager) | 1793

show accounting pending-accounting-stops

IN THIS SECTION

- Syntax | 2669
- Description | 2669
- Options | 2669
- Required Privilege Level | 2670
- Output Fields | 2670
- Sample Output | 2672
- Release Information | 2674

Syntax

```
show accounting pending-accounting-stops
<detail | terse>
<profile-name>
```

Description

Display all statistics for all pending accounting stop requests, including both service and session requests.

Options

none	Display information for all access profiles.
detail | terse	(Optional) Display the specified level of output.
profile-name	(Optional) Particular access profile for which you want to display accounting stop statistics.

Required Privilege Level

view

Output Fields

Table 96 on page 2670 lists the output fields for the `show accounting pending-accounting-stops` command. Output fields are listed in the approximate order in which they appear.

Table 96: show accounting pending-accounting-stops Output Fields

Field Name	Field Description	Level of Output
`Type`	Type of client.	All levels
`Username`	Name of the user logged in to the session.	All levels
`Logical system/Routing instance`	Logical system and routing instance used for the session.	`detail` none
`Access-profile`	Access profile used for AAA services for the session.	`detail` none
`Session ID`	ID of the subscriber session; generated when the subscriber logs in. In the `Service name` block, this is the ID of the service session.	All levels
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	`detail` none
`IP Address`	IP address of the subscriber.	`detail` none
`IPv6 Prefix`	IPv6 address of the subscriber.	`detail` none

Table 96: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Authentication State	State of the subscriber authentication session: AuthInit, AuthStart, AuthChallenge, AuthRedirect, AuthClntRespWait, AuthAcctVolStatsAckWait, AuthAcctStopAckWait, AuthServCreateRespWait, AuthLogoutStart, AuthStateActive, AuthClntLogoutRespWait, AuthProfileUpdateWait, AuthProvisionRespWait, AuthProvisionServiceCreationWait	detail none
Accounting State	State of the subscriber accounting session: Acc-Init, Acc-Start-Sent, Imm-Update-Stats-Pending, Acc-Interim-Sent, Acc-Stop-Stats-Pending, Acc-Stop-Sent, Acc-Stop-On-Fail-Deny-Sent, Acc-Stop-Ackd	detail none
Service name	Name of the attached service or policy.	detail none
Service State	State of the service provided in the subscriber session.	detail none
Session uptime	How long the session has been up, in *HH:MM:SS*.	detail none
Accounting status	Status of the accounting configuration for the service, on or off, and the type of accounting, time or volume+time. Configured in RADIUS Service-Statistics VSA [26-69].	detail none
Service accounting session ID	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the accounting-session-id-format statement.	detail none
Service accounting state	State of the service accounting session: Acc-Init, Acc-Start-Sent, Imm-Update-Stats-Pending, Acc-Interim-Sent, Acc-Stop-Stats-Pending, Acc-Stop-Sent, Acc-Stop-On-Fail-Deny-Sent, Acc-Stop-Ackd	detail none

Table 96: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting interim interval	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail none
Subscriber ID	ID of the subscriber; generated when the subscriber logs in.	detail none
Service ID	ID of the subscriber service.	All levels
Service	Name of the attached service or policy.	terse

Sample Output

show accounting pending-accounting-stops detail

```
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: 2010:db8:9999:18::/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
```

show accounting pending-accounting-stops (Specific 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
```

show accounting pending-accounting-stops terse

```
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
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- request network-access aaa replay pending-accounting-stops | 2641
- show network-access aaa statistics pending-accounting-stops | 2979

show ancp cos

IN THIS SECTION

- Syntax | 2674
- Description | 2674
- Options | 2675
- Required Privilege Level | 2675
- Output Fields | 2675
- Sample Output | 2678
- Release Information | 2681

Syntax

```
show ancp cos
<identifier identifier>
<last-update>
<pending-update>
```

Description

Display information about the CoS state for subscriber traffic.

Options

identifier *identifier* (Optional) Display information about the local loops for the specified access identifier.

last-update (Optional) Display the most recently updated CoS information.

pending-update (Optional) Display the pending update of CoS information.

Required Privilege Level

view

Output Fields

Table 97 on page 2675 lists the output fields for the `show ancp cos` command. Output fields are listed in the approximate order in which they appear.

Table 97: show ancp cos Output Fields

Field Name	Field Description
`Per-DSL CoS adjustment`	Adjustment values applied by the ANCP agent to the actual downstream rates and frame overhead for frame-mode DSL types. The agent then reports the adjusted rates to CoS to establish a shaping rate for the CoS node that corresponds to the subscriber access line.
`QoS Adjust Flag`	State of QoS adjust: • `TRUE`—The ANCP agent is enabled to adjust the actual downstream data rates and frame overhead and report the adjusted values to CoS. • `FALSE`—The ANCP agent is not enabled to adjust and report values to CoS.
`ADSL bytes`	Number of bytes by which the actual ADSL downstream cell overhead is adjusted before reporting it to CoS.
`ADSL2 bytes`	Number of bytes by which the actual ADSL2 downstream cell overhead is adjusted before reporting it to CoS.

Table 97: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
ADSL2-PLUS bytes	Number of bytes by which the actual ADSL2+ downstream cell overhead is adjusted before reporting it to CoS.
SDSL overhead adjusted	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to CoS.
SDSL bytes	Number of bytes by which the actual SDSL downstream frame overhead is adjusted before reporting it to CoS.
OTHER overhead adjusted	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to CoS.
OTHER bytes	Number of bytes by which the actual OTHER downstream frame overhead is adjusted before reporting it to CoS.
VDSL overhead adjusted	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to CoS.
VDSL bytes	Number of bytes by which the actual VDSL downstream frame overhead is adjusted before reporting it to CoS.
VDSL2 overhead adjusted	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to CoS.
VDSL2 bytes	Number of bytes by which the actual VDSL2 downstream frame overhead is adjusted before reporting it to CoS.
Per-DSL adjustment for reporting	Adjustment values applied by the ANCP agent to the actual downstream rates for individual DSL types to account for traffic overhead. The agent then reports the adjusted rates to AAA.
ADSL adjustment factor	Percentage by which the actual ADSL downstream rate is adjusted before reporting it to AAA.

Table 97: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
ADSL2 adjustment factor	Percentage by which the actual ADSL2 downstream rate is adjusted before reporting it to AAA.
ADSL2+ adjustment factor	Percentage by which the actual ADSL2+ downstream rate is adjusted before reporting it to AAA.
VDSL adjustment factor	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to AAA.
VDSL2 adjustment factor	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to AAA.
SDSL adjustment factor	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to AAA.
OTHER adjustment factor	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to AAA.
Keepalive Timer	Interval between the keepalive messages that the ANCP agent sends to CoS.
Cos State	State of the interaction between the ANCP agent and CoS: • ANCPD_COS_CONNECT_NEEDED • ANCPD_COS_CONNECT_PENDING • ANCPD_COS_CONNECT_DONE • ANCPD_COS_SESSION_SENT • ANCPD_COS_WRITE_READY
Connect Time	Time at which the ANCP agent connected to CoS; useful for debugging.

Table 97: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
Session Time	Time at which the ANCP agent sent a session connect message to CoS; useful for debugging.
Routing Instance Time	Time at which the ANCP agent sent the routing instance to CoS; useful for debugging.
Keepalive Time	Time at which the last keepalive message was sent.
Update Time	Time at which the shaping rate was last updated.
Type	Subscriber access type: ifl indicates that a single VLAN carries subscriber traffic and iflset indicates that a set of VLANs carries subscriber traffic.
Name	System-wide name of the particular subscriber access.
Index	Access identifier.
Pending Update	Actual downstream data rate to be applied next to this local loop, in Kbps.
Last Update	Adjusted downstream data rate last reported to CoS by the ANCP agent for this local loop, in Kbps.

Sample Output

show ancp cos

```
user@host> show ancp cos

Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:               20
```

```
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100


Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:14 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name             Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                 36000 Kbps
```

show ancp cos last-update

```
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:              20
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
```

```
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100

Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name             Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                 36000 Kbps
```

show ancp cos pending-update

```
user@host> show ancp cos pending-update
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30
```

```
Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100

Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

- show ancp neighbor | 2681
- show ancp statistics | 2693
- show ancp subscriber | 2699

show ancp neighbor

IN THIS SECTION

- Syntax | 2682
- Description | 2682
- Options | 2682
- Required Privilege Level | 2682

- Output Fields | 2682
- Sample Output | 2688
- Release Information | 2692

Syntax

```
show ancp neighbor
<brief | detail>
<ip-address ip-address
<system-name mac-address>
```

Description

Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.

Options

brief | detail	(Optional) Display the specified level of detail.
ip-address *ip-address*	(Optional) Display information about the neighbor (access node) specified by the IP address.
system-name *mac-address*	(Optional) Display information about the neighbor (access node) specified by the MAC address.

Required Privilege Level

view

Output Fields

Table 98 on page 2683 lists the output fields for the `show ancp neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 98: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: • `0x31`—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*. • `0x32`—ANCP version 1, defined in *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*.	`brief detail` none
`IP Address`	IP address of the ANCP neighbor.	`brief detail` none
`PartId`	Number that associates the ANCP message with a specific partition.	`brief` none
`State`	Operational state of the ANCP adjacency: • Configured—The neighbor has been configured, but has never been in the `Established` state. An asterisk (*) is prefixed to the neighbor entry for this state. • Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly. • `Established`—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established. • Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the `Established` state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state.	All levels
`Time`	How long the adjacency has been up in one of the following formats: • *nwndnh*—number of weeks, days, and hours • *nd hh:mm:ss*—number of days, hours, minutes, and seconds	`brief detail` none

Table 98: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Subscriber Count`	Number of subscribers associated with the ANCP neighbor (access local loop).	`brief` none
`Capabilities`	Negotiated ANCP capability: • `Topo`—Topology discovery. • `OAM`—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router.	All levels
`System Name`	MAC address of the ANCP neighbor.	`detail`
`TCP Port`	TCP port on which ANCP messages are exchanged.	`detail`
`System Instance`	Number identifying the ANCP link instance from the edge device's perspective.	`detail`
`Peer Instance`	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	`detail`
`Timer`	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	`detail`
`Partition Type`	Number that identifies whether partitions are used and how the ID is negotiated: • `0`—No partition. • `1`—Fixed partition requested. • `2`—Fixed partition assigned.	`detail`

Table 98: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Partition Flag	Number that specifies the type of partition requested: 1 (new adjacency) or 2 (recovered adjacency).	detail
Partition Identifier	Number that identifies a logical partition of an access node with which the ANCP agent has formed an adjacency. A value of zero indicates that the 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. A nonzero value indicates that the agent supports each neighbor on an IP address over a single TCP session with a nonzero partition ID.	detail
Partition Adjacencies	Number of adjacencies that share the partition.	detail
Dead Timer	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	detail
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.	detail
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.	detail
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	detail
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	detail

Table 98: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Received Port Up Count`	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	`detail`
`Received Port Down Count`	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	`detail`
`Received Generic Resp Count`	Number of generic response messages received from neighbors.	`detail`
`Received Adjacency Update Count`	Number of adjacency update messages received from neighbors.	`detail`
`Received OAM Count`	Number of OAM responses received from neighbors in reply to `request` commands.	`detail`
`Received Other Count`	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	`detail`
`Sent Syn Count`	Number of synchronization messages sent to neighbors to maintain adjacencies.	`detail`
`Sent Synack Count`	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	`detail`
`Sent Rstack Count`	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	`detail`
`Sent Ack Count`	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	`detail`
`Sent Generic Resp Count`	Number of generic response messages sent to neighbors.	`detail`

Table 98: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sent OAM Count`	Number of OAM `request` commands sent to neighbors.	`detail`
`Max Discovery Limit Exceed Count`	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	`detail`
`Result Codes`	Number of generic response messages sent to neighbors that include each of the following result codes: • `Invalid Request Message Count`—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission. • `Specified Port(s) Down Count`—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • `Out of Resources Count`—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node. • `Request Msg Not Implemented Count`— • `Malformed Msg Count`—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • `TLV Missing Count`—One or more mandatory TLVs was missing from a request. • `Invalid TLV Contents Count`—The contents of one or more TLVs in the request do not match its required specification. • `Non-Existent Port(s) Count`—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.	`detail`

Sample Output

show ancp neighbor

```
user@host> show ancp neighbor
  Version IP Address       PartID  State         Time        Subscriber    Capabilities
                                                             Count
  0x31    203.0.113.13         0       Established   11:24       2             Topo
  0x31    203.0.113.15         0       Not Estblshd  2:45        2             Topo
* 0x0     198.51.100.102       0       Establishing  0           0
* 0x0     192.0.2.0            0       Configured    0           0
* 0x0     192.0.2.1            0       Configured    0           0
```

show ancp neighbor detail

```
user@host> show ancp neighbor detail
  Neighbor Information
      Version                 : 0x31
      IP Address              : 192.0.2.85
      System Name             : 00:00:5e:00:53:01
          Up Time                          : 26
          TCP Port                         : 32666
          State                            : Established
          Subscriber Count                 : 4
          Capabilities                     : Topo
          System Instance                  : 2
          Peer Instance                    : 20
          Adjacency Timer (in 100ms)       : 100
          Peer Adjacency Timer (in 100ms)  : 100
          Partition Type                   : 0
          Partition Flag                   : 1
          Partition Identifier             : 0
          Partition Adjacencies            : 0
          Dead Timer                       : 23
          Received Syn Count               : 1
          Received Synack Count            : 1
          Received Rstack Count            : 0
          Received Ack Count               : 4
          Received Port Up Count           : 10
          Received Port Down Count         : 0
```

```
        Received Generic Resp Count      : 0
        Received Adjacency Update Count  : 0
        Received OAM Count               : 0
        Received Other Count             : 0
        Sent Syn Count                   : 1
        Sent Synack Count                : 2
        Sent Rstack Count                : 0
        Sent Ack Count                   : 3
        Sent Generic Resp Count          : 0
        Sent OAM Count                   : 0
        Max Discovery Limit Exceed Count : 0
    Result Codes:                          Received              Sent
        Invalid Request Message Count    : 0                     0
        Specified Port(s) Down Count     : 0                     0
        Out of Resources Count           : 0                     0
        Request Msg Not Implemented Count: 0                     0
        Malformed Msg Count              : 0                     0
        TLV Missing Count                : 0                     0
        Invalid TLV Contents Count       : 0                     0
        Non-Existent Port(s) Count       : 0                     0


    Version                 : 0x32
    IP Address              : 192.168.9.1
    System Name             : 00:00:5e:00:53:02
        Up Time                          : 36
        TCP Port                         : 61408
        State                            : Not Established
        Subscriber Count                 : 1
        Capabilities                     : Topology Discovery
        System Instance                  : 12
        Peer Instance                    : 1
        Adjacency Timer (in 100ms)       : 50
        Peer Adjacency Timer (in 100ms)  : 100
        Partition Type                   : 0
        Partition Flag                   : 1
        Partition Identifier             : 0
        Partition Adjacencies            : 0
        Dead Timer                       : 23
        Received Syn Count               : 24
        Received Synack Count            : 20
        Received Rstack Count            : 2
        Received Ack Count               : 9
```

```
            Received Port Up Count           : 5
            Received Port Down Count         : 0
            Received Generic Resp Count      : 0
            Received Adjacency Update Count  : 0
            Received OAM Responses Count     : 2
            Received Other Count             : 0
            Sent Syn Count                   : 20
            Sent Synack Count                : 24
            Sent Rstack Count                : 1
            Sent Generic Resp Count          : 0
            Sent Ack Count                   : 9
            Sent OAM Requests Count          : 4
            Max Discovery Limit Exceed Count : 0
        Result Codes:                          Received              Sent
            Invalid Request Message Count    : 0                     0
            Specified Port(s) Down Count     : 0                     0
            Out of Resources Count           : 0                     0
            Request Msg Not Implemented Count: 0                     0
            Malformed Msg Count              : 0                     0
            TLV Missing Count                : 0                     0
            Invalid TLV Contents Count       : 0                     0
            Non-Existent Port(s) Count       : 0                     0
```

show ancp neighbor ip-address

```
user@host> show ancp neighbor ip-address 192.0.2.85

Neighbor Information
      Version                 : 0x32
      IP Address   : 192.0.2.85
      System Name  : 00:00:5e:00:53:ba
          Up Time                          : 26
          TCP Port                         : 32666
          State                            : Established
          Subscriber Count                 : 4
          Capabilities                     : Topo
          System Instance                  : 2
          Peer Instance                    : 20
          Adjacency Timer (in 100ms)       : 100
          Peer Adjacency Timer (in 100ms)  : 100
          Partition Type                   : 0
```

```
          Partition Flag                  : 1
          Partition Identifier            : 0
          Partition Adjacencies           : 0
          Dead Timer                      : 23
          Received Syn Count              : 1
          Received Synack Count           : 1
          Received Rstack Count           : 0
          Received Ack Count              : 4
          Received Port Up Count          : 10
          Received Port Down Count        : 0
          Received Generic Resp Count     : 0
          Received Adjacency Update Count  : 0
          Received OAM Count              : 0
          Received Other Count            : 0
          Sent Syn Count                  : 1
          Sent Synack Count               : 2
          Sent Rstack Count               : 0
          Sent Ack Count                  : 3
          Sent Generic Resp Count         : 0
          Sent OAM Count                  : 0
          Max Discovery Limit Exceed Count : 0
      Result Codes:                         Received              Sent
          Invalid Request Message Count   : 0                     0
          Specified Port(s) Down Count    : 0                     0
          Out of Resources Count          : 0                     0
          Request Msg Not Implemented Count: 0                    0
          Malformed Msg Count             : 0                     0
          TLV Missing Count               : 0                     0
          Invalid TLV Contents Count      : 0                     0
          Non-Existent Port(s) Count      : 0                     0
```

show ancp neighbor system-name

```
user@host> show ancp neighbor 00:00:5e:00:53:ba detail

Neighbor Information
   Version                : 0x31
   IP Address             : 203.0.113.101
   System Name            : 00:00:5e:00:53:ba
     Up Time                           : 19
     TCP Port                          : 1028
```

```
    State                               : Established
    Subscriber Count                    : 2
    Capabilities                        : Topology Discovery, OAM
    System Instance                     : 1
    Peer Instance                       : 10
    Adjacency Timer (in 100ms)          : 100
    Peer Adjacency Timer (in 100ms)     : 250
    Partition Type                      : 0
    Partition Flag                      : 1
    Partition Identifier                : 0
    Partition Adjacencies               : 0
    Dead Timer                          : 55
    Received Syn Count                  : 1
    Received Synack Count               : 1
    Received Rstack Count               : 0
    Received Ack Count                  : 1
    Received Port Up Count              : 34
    Received Port Down Count            : 0
    Received Generic Resp Count         : 0
    Received Adjacency Update Count     : 0
    Received OAM Responses Count        : 2
    Received Other Count                : 0
    Sent Syn Count                      : 1
    Sent Synack Count                   : 1
    Sent Rstack Count                   : 0
    Sent Ack Count                      : 3
    Sent Generic Resp Count             : 0
    Sent OAM Requests Count             : 4
    Max Discovery Limit Exceed Count    : 3
  Result Codes:                         Received              Sent
    Invalid Request Message Count       : 0                     0
    Specified Port(s) Down Count        : 0                     0
    Out of Resources Count              : 0                     0
    Request Msg Not Implemented Count   : 0                     0
    Malformed Msg Count                 : 0                     0
    TLV Missing Count                   : 0                     0
    Invalid TLV Contents Count          : 0                     0
    Non-Existent Port(s) Count          : 0                     0
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ancp cos | 2674

show ancp subscriber | 2699

show ancp statistics

IN THIS SECTION

- Syntax | 2693
- Description | 2693
- Options | 2693
- Required Privilege Level | 2694
- Output Fields | 2694
- Sample Output | 2696
- Release Information | 2699

Syntax

```
show ancp statistics
<ip-address ip-address>
<system-name mac-address>
```

Description

Display statistics for all ANCP neighbors (access nodes) or the specified ANCP neighbor.

Options

none	Display statistics for all ANCP neighbors, including global statistics not show for individual neighbors.
ip-address *ip-address*	(Optional) Display statistics for only the neighbor with the specified IP address.

system-name *mac-address*	(Optional) Display statistics for only the neighbor with the specified MAC address.

Required Privilege Level

view

Output Fields

Table 99 on page 2694 lists the output fields for the `show ancp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 99: show ancp statistics Output Fields

Field Name	Field Description
`Number of neighbors`	Total count of ANCP neighbors.
`Number of subscribers`	Total count of ANCP subscribers.
`Accept Count`	Number of neighbor TCP/IP sessions accepted on listener socket.
`Accept Fail Count`	Number of neighbor TCP/IP sessions that failed due to one of the following causes: session already exists, maximum number of ANCP connections exceeded, creation of session or neighbor failed, or protocol start failed.
`No Config Accept Deny Count`	Number of neighbor TCP/IP sessions that failed because the neighbor was not configured.
`Received Syn Count`	Number of synchronization messages received from neighbors to maintain adjacencies.
`Received Synack Count`	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.
`Received Rstack Count`	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.

Table 99: show ancp statistics Output Fields *(Continued)*

Field Name	Field Description
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.
Received Generic Resp Count	Number of generic response messages received from neighbors.
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.
Received OAM Count	Number of OAM responses received from neighbors in reply to request commands.
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.

Table 99: show ancp statistics Output Fields *(Continued)*

Field Name	Field Description
Sent Generic Resp Count	Number of generic response messages sent to neighbors.
Sent OAM Count	Number of OAM request commands sent to neighbors.
Result Codes	Number of generic response messages sent to neighbors that include each of the following result codes: • Invalid Request Message Count—A properly formed request messages violated the protocol because of timing (such as a race condition) or direction of transmission. • Specified Port(s) Down Count—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • Out of Resources Count—the ANCP agent is out of resources, probably not related to the access lines. This result code is sent only by an access node. • Request Msg Not Implemented Count— • Malformed Msg Count—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • TLV Missing Count—One or more mandatory TLVs was missing from a request. • Invalid TLV Contents Count—The contents of one or more TLVs in the request do not match its required specification. • Non-Existent Port(s) Count—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.

Sample Output

show ancp statistics

```
user@host> show ancp statistics
  Statistics
          Number of neighbors             : 4
```

```
          Number of subscribers           : 6
          Accept Count                    : 0
          Accept Fail Count               : 0
          No Config Accept Deny Count     : 0
          Received Syn Count              : 2
          Received Synack Count           : 1
          Received Rstack Count           : 0
          Received Ack Count              : 8
          Received Port Up Count          : 7
          Received Port Down Count        : 0
          Received Generic Resp Count     : 0
          Received Adjacency Update Count : 0
          Received OAM Count              : 0
          Received Other Count            : 0
          Sent Syn Count                  : 1
          Sent Synack Count               : 1
          Sent Rstack Count               : 0
          Sent Ack Count                  : 17
          Sent Generic Resp Count         : 0
          Sent OAM Count                  : 4
      Result Codes:                         Received              Sent
          Invalid Request Message Count    : 0                    0
          Specified Port(s) Down Count     : 0                    0
          Out of Resources Count           : 0                    0
          Request Msg Not Implemented Count: 0                    0
          Malformed Msg Count              : 0                    0
          TLV Missing Count                : 0                    0
          Invalid TLV Contents Count       : 0                    0
          Non-Existent Port(s) Count       : 0                    0
```

show ancp statistics ip-address

```
user@host> show ancp statistics ip-address 203.0.113.1
  Statistics
          Received Syn Count              : 2
          Received Synack Count           : 1
          Received Rstack Count           : 0
          Received Ack Count              : 8
          Received Port Up Count          : 7
          Received Port Down Count        : 0
```

```
          Received Generic Resp Count     : 0
          Received Adjacency Update Count : 0
          Received OAM Count              : 0
          Received Other Count            : 0
          Sent Syn Count                  : 1
          Sent Synack Count               : 1
          Sent Rstack Count               : 0
          Sent Ack Count                  : 17
          Sent Generic Resp Count         : 0
          Sent OAM Count                  : 4
      Result Codes:                          Received              Sent
          Invalid Request Message Count    : 0                     0
          Specified Port(s) Down Count     : 0                     0
          Out of Resources Count           : 0                     0
          Request Msg Not Implemented Count: 0                     0
          Malformed Msg Count              : 0                     0
          TLV Missing Count                : 0                     0
          Invalid TLV Contents Count       : 0                     0
          Non-Existent Port(s) Count       : 0                     0
```

show ancp statistics system-name

```
user@host> show ancp statistics system-name 00:00:5E:00:53:02
  Statistics
          Received Syn Count              : 2
          Received Synack Count           : 1
          Received Rstack Count           : 0
          Received Ack Count              : 8
          Received Port Up Count          : 7
          Received Port Down Count        : 0
          Received Generic Resp Count     : 0
          Received Adjacency Update Count : 0
          Received OAM Count              : 0
          Received Other Count            : 0
          Sent Syn Count                  : 1
          Sent Synack Count               : 1
          Sent Rstack Count               : 0
          Sent Ack Count                  : 17
          Sent Generic Resp Count         : 0
          Sent OAM Count                  : 4
```

```
    Result Codes:                          Received            Sent
        Invalid Request Message Count    : 0                   0
        Specified Port(s) Down Count     : 0                   0
        Out of Resources Count           : 0                   0
        Request Msg Not Implemented Count: 0                   0
        Malformed Msg Count              : 0                   0
        TLV Missing Count                : 0                   0
        Invalid TLV Contents Count       : 0                   0
        Non-Existent Port(s) Count       : 0                   0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- show ancp cos | 2674
- show ancp neighbor | 2681
- show ancp subscriber | 2699

show ancp subscriber

IN THIS SECTION

- Syntax | 2700
- Description | 2700
- Options | 2700
- Required Privilege Level | 2701
- Output Fields | 2701
- Sample Output | 2705
- Release Information | 2708

Syntax

```
show ancp subscriber
<brief | detail>
<access-aggregation-circuit-id circuit-identifier>
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Description

Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACI, or the subscriber associated with the specified ANCP neighbor (access node).

After an ancpd restart, this command displays orphaned entries (marked with an o) for subscriber sessions that were established before the restart but which have not yet been reestablished. As sessions are reestablished, the number of orphaned entries displayed by the command decreases. The number reaches zero when all sessions are reestablished or when the orphaned-interface timer expires.

Options

none	Display information about all subscribers.
brief | detail	(Optional) Display the specified level of detail.
access-aggregation-circuit-id *circuit-identifier*	(Optional) Display information about ANCP subscribers whose Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) matches the specified value. A *circuit-identifier* that begins with the# character indicates a backhaul line identifier. You can specify a wildcard (*) anywhere in the string.
identifier *identifier*	(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.
ip-address *ip-address*	(Optional) Display information about the subscribers connected to the access node specified by the IP address.
system-name *mac-address*	(Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

Output Fields

Table 100 on page 2701 lists the output fields for the `show ancp subscriber` command. Output fields are listed in the approximate order in which they appear.

Table 100: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
`Loop Identifier`	Access loop identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established. An `o` indicates that the entry is for an orphaned interface and represents a previously established subscriber session that has not been reestablished after an ancpd restart. The number of orphaned entries decreases as the ANCP neighbors reestablish adjacencies and the protocol subscriber sessions are reestablished. The command output indicates this by removing the `o` marker. Eventually the number of orphaned entries reaches zero, because either all the adjacencies and subscriber sessions have been reestablished or any remaining orphaned entries are removed when the orphaned-interface timer expires.	`brief` none
`DSL Line State`	State of the DSL line: `Idle`, `Showtime`, or `Silent`.	`brief detail`
`Access Type`	Type of access line employed by the access node: `ADSL1`, `ADSL2`, `ADSL2+`, `VDSL1`, `VDSL2`, `SDSL`, `G.fast`, `VDSL2 Annex Q`, `SDSL bonded`, `VDSL2 bonded`, `G.fast bonded` `VDSL2 Annex Q bonded` or `OTHER`.	`brief detail` none

Table 100: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface`	Name of the interface set or logical interface.	`brief detail` none
`Rate Kbps`	Actual downstream data rate for this local loop.	`brief` none
`Neighbor`	IP address of ANCP neighbor (access node).	`brief` none
`Access Loop Circuit Identifier`	Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.	`detail`
`Neighbor IP Address`	IP address of the ANCP neighbor (access node).	`detail`
`Aggregate Circuit Identifier`	ASCII identifier for the subscriber access loop; value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003).	`detail`
`Aggregate Circuit Identifier Binary`	Binary identifier for the VLAN circuit ID.	`detail`
`Tech Type`	Type of technology employed by the subscriber. Currently Junos OS supports `DSL` technology type only.	`detail`
`DSL Line Data Link`	Data link protocol employed on the access loop: `AAL5` or `Ethernet`.	`detail`

Table 100: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
DSL Line Encapsulation	Encapsulation type on the access loop, for Ethernet only: • 0—NA, type not conveyed • 1—Untagged Ethernet • 2—Single-tagged Ethernet	detail
DSL Line Encapsulation Payload	Payload carried across the access loop: • 0—NA, type not conveyed • 1—PPPoA LLC • 2—PPPoA null • 3—IPoA LLC • 4—IPoA null • 5—Ethernet over AAL5 LLC with FCS • 6—Ethernet over AAL5 LLC without FCS • 7—Ethernet over AAL5 null with FCS • 8—Ethernet over AAL5 null without FCS	detail
Interface Type	Type of interface employed for subscriber traffic: ifl for a single VLAN or interface-set for a configured group of VLANs.	detail
Actual Net Data Upstream	Actual upstream data rate for this local loop, in Kbps.	detail
Actual Net Data Downstream	Actual downstream data rate for this local loop, in Kbps.	detail

Table 100: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Minimum Net Data Upstream	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Minimum Net Data Downstream	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Upstream	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Downstream	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Attainable Net Data Upstream	Maximum attainable upstream data rate for this local loop, in Kbps.	detail
Attainable Net Data Downstream	Maximum attainable downstream data rate for this local loop, in Kbps.	detail
Minimum Low Power Data Downstream	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Minimum Low Power Data Upstream	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Maximum Interleave Delay Downstream	Maximum interleaving delay for downstream data, in milliseconds.	detail
Maximum Interleave Delay Upstream	Maximum interleaving delay for upstream data, in milliseconds.	detail
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	detail

Table 100: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	detail

Sample Output

show ancp subscriber

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Tech Type     Access Type   Interface       Rate   Neighbor
                         State                                 Kbps
**circuit 101            Idle      DSL       ADSL1    ----             32       203.0.113.13
**circuit 102            Idle      DSL       ADSL1    ----             32       203.0.113.13
  circuit 301            Showtime  DSL       ADSL1    ----             32       203.0.113.15
  circuit 302            Showtime  DSL       ADSL1    ----             32       203.0.113.15
```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Tech Type     Access Type  Interface         Rate     Neighbor
                         State                                 Kbps
o circuit 201            Showtime  DSL        ADSL1    ----             222222
o circuit 202            Showtime  DSL        ADSL1    ----             222222
```

show ancp subscriber brief

```
user@host> show ancp subscriber brief

  Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10                VDSL2     set-ge-10410      64       203.0.113.102
  port-1-11                VDSL2     set-ge-10411      64       203.0.113.111
```

```
  port-2-10               VDSL2    ge-1/0/4.12        64       203.0.113.112
  port-2-11               VDSL2    ge-1/0/4.13        64       203.0.113.113
```

show ancp subscriber detail

```
user@host> show ancp subscriber detail
 Subscriber Information
*  Access Loop Circuit Identifier : circuit 101
     Neighbor IP Address                  : 203.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                                                                : DSL
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
*  Access Loop Circuit Identifier: circuit 102
     Neighbor IP Address                  : 213.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                                                                : DSL
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
```

```
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
...
```

show ancp subscriber access-aggregation-circuit-id detail

```
user@host> show ancp subscriber access-aggregation-circuit-id "#TEST-DPU-C-100" detail
Subscriber Information
 * Access Loop Circuit Identifier : circuit 201
     Neighbor IP Address                  : 192.0.2.1
     Access Loop Remote Identifier        : remote 123
     Aggregate Circuit Identifier         : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                           : DSL
     Interface Type                       : interface
     Interface                            : ge-1/0/0.3221225475
     Actual Net Data Upstream             : 1024
     Actual Net Data Downstream           : 2048
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
* Access Loop Circuit Identifier : circuit 202
     Neighbor IP Address                  : 192.0.2.1
     Access Loop Remote Identifier        : remote 185
     Aggregate Circuit Identifier         : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                           : DSL
     Interface Type                       : interface
```

```
    Interface                          : ge-1/0/0.3221225476
    Actual Net Data Upstream           : 1024
    Actual Net Data Downstream         : 2048
    Maximum Net Data Upstream          : 0
    Maximum Net Data Downstream        : 0
```

show ancp subscriber identifier identifier-string detail

```
user@host> show ancp subscriber identifier port-1-11 detail


     Access Loop Identifier : port-1-11
         Neighbor IP Address                 : 203.0.113.112
         Aggregate Circuit Identifier Binary  : 0/0
         DSL Type                            :  DSL  0
         Interface Type                      : interface-set
         Interface                           : set-ge-10411
         DSL Line State                      : Show Time
         Actual Net Data Upstream            : 64
         Actual Net Data Downstream          : 64
         DSL Line Data Link                  : AAL5
         DSL Line Encapsulation              : N/A
         DSL Line Encapsulation Payload      : N/A
         Minimum Net Data Upstream           : 64
         Minimum Net Data Downstream         : 64
         Maximum Net Data Upstream           : 64
         Maximum Net Data Downstream         : 64
         Attainable Net Data Upstream        : 64
         Attainable Net Data Downstream      : 64
         Minimum Low Power Data Downstream   : 64
         Minimum Low Power Data Upstream     : 64
         Maximum Interleave Delay Downstream  : 50
         Maximum Interleave Delay Upstream   : 50
         Actual Interleave Delay Downstream  : 50
         Actual Interleave Delay Upstream    : 50
```

Release Information

Command introduced in Junos OS Release 9.4.

`neighbor` option replaced with `ip-address` in Junos OS Release 16.1.

`system-name` option introduced in Junos OS Release 16.1.

`access-aggregation-circuit-id` option introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

clear ancp subscriber | 2547

show ancp cos | 2674

show ancp neighbor

show ancp statistics | 2693

show ancp summary

IN THIS SECTION

- Syntax | 2709
- Description | 2709
- Options | 2710
- Required Privilege Level | 2710
- Output Fields | 2710
- Sample Output | 2711
- Release Information | 2711

Syntax

```
show ancp summary
```

Description

Display a summary of the counts and states for all ANCP neighbors and subscribers.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 101 on page 2710 lists the output fields for the `show ancp summary` command. Output fields are listed in the approximate order in which they appear.

Table 101: show ancp summary Output Fields

Field Name	Field Description
`Configured`	Number of ANCP neighbors in the `Configured` state; that is, that have been configured but never established.
`Establishing`	Number of ANCP neighbors in the `Establishing` state; that is, where negotiations are in progress.
`Established`	Number of ANCP neighbors in the `Established` state; that is, where negotiations have succeeded and the ANCP session has been established.
`Not Estblshd`	Number of ANCP neighbors in the `Not Estblshd` state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
`Total`	Total number of ANCP neighbors; sum of neighbors in the `Configured`, `Establishing`, `Established`, and `Not Estblshd` states.
`Showtime`	Number of DSL lines in `Showtime` state.
`Idle`	Number of DSL lines in `Idle` state.
`Silent`	Number of DSL lines in `Silent` state.

Table 101: show ancp summary Output Fields *(Continued)*

Field Name	Field Description
Unknown	Number of DSL lines where the state is not Showtime, Idle, or Silent.
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the Showtime, Idle, Silent, and Unknown states.

Sample Output

show ancp summary

```
user@host> show ancp summary

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24

Subscribers Summary:
  Showtime      Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           4          0           0          0          4
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- show ancp neighbor | 2681
- show ancp summary neighbor | 2712
- show ancp subscriber | 2699
- show ancp summary subscriber | 2715

show ancp summary neighbor

IN THIS SECTION

- Syntax | 2712
- Description | 2712
- Options | 2712
- Required Privilege Level | 2712
- Output Fields | 2713
- Sample Output | 2714
- Release Information | 2714

Syntax

```
show ancp summary neighbor
<ip-address ip-address | system-name mac-address>
```

Description

Display a summary of the counts and states for all ANCP neighbors and of the neighbor's subscribers when you specify a particular neighbor.

Options

ip-address *ip-address*	(Optional) IP address of the ANCP neighbor (access node).
system-name *mac-address*	(Optional) MAC address of the ANCP neighbor (access node).

Required Privilege Level

view

Output Fields

Table 102 on page 2713 lists the output fields for the `show ancp summary` command. Output fields are listed in the approximate order in which they appear.

Table 102: show ancp summary neighbor Output Fields

Field Name	Field Description
`Configured`	Number of ANCP neighbors in the `Configured` state; that is, that have been configured but never established.
`Establishing`	Number of ANCP neighbors in the `Establishing` state; that is, where negotiations are in progress.
`Established`	Number of ANCP neighbors in the `Established` state; that is, where negotiations have succeeded and the ANCP session has been established.
`Not Estblshd`	Number of ANCP neighbors in the `Not Estblshd` state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
`Total`	Total number of ANCP neighbors; sum of neighbors in the `Configured`, `Establishing`, `Established`, and `Not Estblshd` states.
`Showtime`	Number of DSL lines for the neighbor in `Showtime` state.
`Idle`	Number of DSL lines for the neighbor in `Idle` state.
`Silent`	Number of DSL lines for the neighbor in `Silent` state.
`Unknown`	Number of DSL lines for the neighbor where the state is not `Showtime`, `Idle`, or `Silent`.
`Total`	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the `Showtime`, `Idle`, `Silent`, and `Unknown` states.

Sample Output

show ancp summary neighbor

```
user@host> show ancp summary neighbor

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24
```

show ancp summary neighbor (IP Address)

```
user@host> show ancp summary neighbor ip-address 192.168.10.1

Neighbor Summary:192.168.10.1 status Established

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           6          0           0          0          6
```

show ancp summary neighbor (MAC Address)

```
user@host> show ancp summary neighbor system-name 00:00:5E:00:53:02

Neighbor Summary:00:00:5E:00:53:02 status Established

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           5          1           2          0          8
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

show ancp summary | 2709

show ancp subscriber | 2699

show ancp summary subscriber | 2715

show ancp summary subscriber

IN THIS SECTION

- Syntax | 2715
- Description | 2715
- Options | 2715
- Required Privilege Level | 2715
- Output Fields | 2716
- Sample Output | 2716
- Release Information | 2716

Syntax

```
show ancp summary subscriber
```

Description

Display a summary of the counts and states for all ANCP subscribers.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 103 on page 2716 lists the output fields for the `show ancp summary subscriber` command. Output fields are listed in the approximate order in which they appear.

Table 103: show ancp summary subscriber Output Fields

Field Name	Field Description
`Showtime`	Number of DSL lines in `Showtime` state.
`Idle`	Number of DSL lines in `Idle` state.
`Silent`	Number of DSL lines in `Silent` state.
`Unknown`	Number of DSL lines where the state is not `Showtime`, `Idle`, or `Silent`.
`Total`	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the `Showtime`, `Idle`, `Silent`, and `Unknown` states.

Sample Output

show ancp summary subscriber

```
user@host> show ancp summary subscriber

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           8          1           0          1         10
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

show ancp summary | 2709

show ancp neighbor | 2681

show ancp summary neighbor | 2712

show class-of-service interface

IN THIS SECTION

- Syntax | 2717
- Description | 2717
- Options | 2718
- Required Privilege Level | 2718
- Output Fields | 2719
- Sample Output | 2739
- Release Information | 2761

Syntax

```
show class-of-service interface <comprehensive | detail> <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

NOTE: On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

```
error: the class-of-service subsystem is not running
```

Options

none Display CoS associations for all physical and logical interfaces.

comprehensive (M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

detail (M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the `interface` *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the `interface` *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

interface-name (Optional) Display class-of-service (CoS) associations for the specified interface.

none Display CoS associations for all physical and logical interfaces.

NOTE: ACX5000 routers do not support classification on logical interfaces and therefore do not show CoS associations for logical interfaces with this command.

Required Privilege Level

view

Output Fields

Table 104 on page 2719 describes the output fields for the `show class-of-service interface` command. Output fields are listed in the approximate order in which they appear.

Table 104: show class-of-service interface Output Fields

Field Name	Field Description
`Physical interface`	Name of a physical interface.
`Index`	Index of this interface or the internal index of this object. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management.
`Dedicated Queues`	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX-Series routers) This field is not displayed for enhanced subscriber management.
`Maximum usable queues`	Number of queues you can configure on the interface.
`Maximum usable queues`	Maximum number of queues you can use.
`Total non-default queues created`	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.
`Rewrite Input IEEE Code-point`	(QFX3500 switches only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the Shaping rate field is displayed for either the physical interface or the logical interface.
Scheduler map	Name of the output scheduler map associated with this interface. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.
Scheduler map forwarding class sets	(QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Traffic-control-profile	Name of the associated traffic control profile. (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1006) instead of with a subscriber interface.
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Congestion-notification	(QFX Series and EX4600 switches only) Congestion notification state, enabled or disabled.
Monitoring Profile Name	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the intrface.
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier, Fragmentation-map (for LSQ interfaces only), Scheduler-map, Rewrite, Translation Table (for IQE PICs only), or traffic-class-map (for T4000 routers with Type 5 FPCs).
Name	Name of an object.
Type	Type of an object: dscp, dscp-ipv6, exp, ieee-802.1, ip, inet-precedence, or ieee-802.1ad (for traffic class map on T4000 routers with Type 5 FPCs)..
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status. • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.
`Device flags`	The `Device flags` field provides information about the physical device and displays one or more of the following values: • `Down`—Device has been administratively disabled. • `Hear-Own-Xmit`—Device receives its own transmissions. • `Link-Layer-Down`—The link-layer protocol has failed to connect with the remote endpoint. • `Loopback`—Device is in physical loopback. • `Loop-Detected`—The link layer has received frames that it sent, thereby detecting a physical loopback. • `No-Carrier`—On media that support carrier recognition, no carrier is currently detected. • `No-Multicast`—Device does not support multicast traffic. • `Present`—Device is physically present and recognized. • `Promiscuous`—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • `Quench`—Transmission on the device is quenched because the output buffer is overflowing. • `Recv-All-Multicasts`—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • `Running`—Device is active and enabled.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Interface flags`	The `Interface flags` field provides information about the physical interface and displays one or more of the following values:
	• `Admin-Test`—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
	• `Disabled`—Interface is administratively disabled.
	• `Down`—A hardware failure has occurred.
	• `Hardware-Down`—Interface is nonfunctional or incorrectly connected.
	• `Link-Layer-Down`—Interface keepalives have indicated that the link is incomplete.
	• `No-Multicast`—Interface does not support multicast traffic.
	• `No-receive No-transmit`—Passive monitor mode is configured on the interface.
	• `Point-To-Point`—Interface is point-to-point.
	• `Pop all MPLS labels from packets of depth`—MPLS labels are removed as packets arrive on an interface that has the `pop-all-labels` statement configured. The depth value can be one of the following:
	  • `1`—Takes effect for incoming packets with one label only.
	  • `2`—Takes effect for incoming packets with two labels only.
	  • `[ 1 2 ]`—Takes effect for incoming packets with either one or two labels.
	• `Promiscuous`—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
	• `Recv-All-Multicasts`—Interface is in multicast promiscuous mode and provides no multicast filtering.
	• `SNMP-Traps`—SNMP trap notifications are enabled.
	• `Up`—Interface is enabled and operational.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Flags	The Logical interface flags field provides information about the logical interface and displays one or more of the following values: • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Link flags`	Provides information about the physical link and displays one or more of the following values: • `ACFC`—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • `Give-Up`—Link protocol does not continue connection attempts after repeated failures. • `Loose-LCP`—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • `Loose-LMI`—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • `Loose-NCP`—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • `Keepalives`—Link protocol keepalives are enabled. • `No-Keepalives`—Link protocol keepalives are disabled. • `PFC`—Protocol field compression is configured. The PPP session negotiates the PFC option.
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.
`CoS queues`	Number of CoS queues configured.
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)`. For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Statistics last cleared`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.
Exclude Overhead Bytes	Exclude the counting of overhead bytes from aggregate queue statistics. • `Disabled`—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • `Enabled`—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • `Enabled for hierarchy`—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Input errors`	Input errors on the interface. The labels are explained in the following list: • `Errors`—Sum of the incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Bucket Drops`—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `HS link FIFO overflows`—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Output errors`	Output errors on the interface. The labels are explained in the following list: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Drops` field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `HS link FIFO underflows`—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeds the MTU of the interface.
`Egress queues`	Total number of egress Maximum usable queues on the specified interface.
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Dropped packets` field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET alarms SONET defects	(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: SONET PHY, SONET section, SONET line, and SONET path.
SONET PHY	Counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET PHY field has the following subfields: • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET section	Counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET section field has the following subfields: • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET line	Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET line field has the following subfields: • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET path field has the following subfields: • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Received SONET overhead`
 	
`Transmitted SONET overhead`
	Values of the received and transmitted SONET overhead:
	• `C2`—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.
	• `F1`—Section user channel byte. This byte is set aside for the purposes of users.
	• `K1` and `K2`—These bytes are allocated for APS signaling for the protection of the multiplex section.
	• `J0`—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.
	• `S1`—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-*N* signal.
	• `Z3` and `Z4`—Allocated for future use.

`Received path trace`
 	
`Transmitted path trace`
	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.

`HDLC configuration`	Information about the HDLC configuration. • `Policing bucket`—Configured state of the receiving policer. • `Shaping bucket`—Configured state of the transmitting shaper. • `Giant threshold`—Giant threshold programmed into the hardware. • `Runt threshold`—Runt threshold programmed into the hardware.
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number. • `PLP byte`—Packet Level Protocol byte.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
CoS information	Information about the CoS queue for the physical interface. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Transmitted Packets`	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values.
`Transmitted Bytes`	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
`Tail-dropped packets`	Number of packets dropped because of tail drop.

`RED-dropped packets`	Number of packets dropped because of random early detection (RED).
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP packets dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP packets dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP packets dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP packets dropped because of RED.
	• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories:
	  • `Low`—Number of low-loss priority packets dropped because of RED.
	  • `Medium-low`—Number of medium-low loss priority packets dropped because of RED.
	  • `Medium-high`—Number of medium-high loss priority packets dropped because of RED.
	  • `High`—Number of high-loss priority packets dropped because of RED.
	NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`RED-dropped bytes`	Number of bytes dropped because of RED. The byte counts vary by PIC type.
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP bytes dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP bytes dropped because of RED.
	NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

`Transmit rate`	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are : • `None`—No rate limit. • `exact`—Queue transmits at the configured rate.
`Buffer size`	Delay buffer size in the queue.
`Priority`	Scheduling priority configured as `low` or `high`.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler: `low`, `medium-low`, `medium-high`, `high`, or `none`.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler.
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.

Table 104: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

Adjustment information	Display the assignment of shaping-rate adjustments on a scheduler node or queue.
	• Adjusting application—Application that is performing the shaping-rate adjustment.
	  • The adjusting application can appear as ancp LS-0, which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes.
	  • The adjusting application can appear as DHCP, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in DHCP option 82, suboption 9 (Vendor Specific Information). The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	  • The adjusting application can also appear as pppoe, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	• Adjustment type—Type of adjustment: absolute or delta.
	• Configured shaping rate—Shaping rate configured for the scheduler node or queue.
	• Adjustment value—Value of adjusted shaping rate.
	• Adjustment target—Level of shaping-rate adjustment performed: node or queue.
	• Adjustment overhead-accounting mode—Configured shaping mode: frame or cell.
	• Adjustment overhead bytes—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.
	• Adjustment target—Level of shaping-rate adjustment performed: node or queue.
	• Adjustment multicast index—

Sample Output

show class-of-service interface (Physical)

```
user@host> show class-of-service interface et-1/0/4
Physical interface: et-1/0/4, Index: 1098
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Monitoring Profile Name: XYZ

  Logical interface: et-1/0/4.16386, Index: 1057
```

show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                  Name                  Type             Index
    Scheduler-map           <default>                               27
    Rewrite                 exp-default           exp               21
    Classifier              exp-default           exp               5
    Classifier              ipprec-compatibility  ip                8
    Forwarding-class-map    exp-default           exp               5
```

show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show class-of-service interface (ANCP)

```
user@host> show class-of-service interface pp0.1073741842
  Logical interface: pp0.1073741842, Index: 341
Object                 Name                   Type                   Index
Traffic-control-profile TCP-CVLAN             Output                 12408
Classifier             dscp-ipv6-compatibility dscp-ipv6                9
Classifier             ipprec-compatibility   ip                       13

    Adjusting application: ancp LS-0
      Adjustment type: absolute
      Configured shaping rate: 4000000
      Adjustment value: 11228000
      Adjustment overhead-accounting mode: Frame Mode
      Adjustment overhead bytes: 50
      Adjustment target: node
```

show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
    Object                 Name                   Type           Index
    Traffic-control-profile tcp-pppoe.o.pp0.1     Output         2726446535
    Classifier             ipprec-compatibility   ip             13

    Adjusting application: PPPoE
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface (DHCP Interface)

```
user@host> show class-of-service interface demux0.1
Logical interface: pp0.1, Index: 85
    Object                 Name                   Type           Index
    Traffic-control-profile tcp-dhcp.o.demux0.1    Output         2726446535
    Classifier             ipprec-compatibility   ip             13
```

```
    Adjusting application: DHCP
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface (T4000 Routers with Type 5 FPCs)

```
user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
    Maximum usable queues: 8, Queues in use: 4
      Shaping rate: 5000000000 bps
      Scheduler map: <default>, Index: 2
      Congestion-notification: Disabled

      Logical interface: xe-4/0/0.0, Index: 77
        Object                  Name                    Type                    Index
        Classifier              ipprec-compatibility    ip                         13
```

show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Physical interface: ge-0/3/0, Index: 138
  Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps
  Scheduler map: interface-schedular-map, Index: 58414
  Input shaping rate: 10000 bps
  Input scheduler map: schedular-map, Index: 15103
  Chassis scheduler map: <default-chassis>, Index: 4
  Congestion-notification: Disabled

  Logical interface ge-0/3/0.0
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
```

```
    inet
    mpls
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet
                           mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up   inet
                           mpls

  Logical interface: ge-0/3/0.0, Index: 68
    Object                  Name                    Type               Index
    Rewrite                 exp-default             exp (mpls-any)       33
    Classifier              exp-default             exp                  10
    Classifier              ipprec-compatibility    ip                   13


  Logical interface ge-0/3/0.1
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    inet
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   inet
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up   inet

  Logical interface: ge-0/3/0.1, Index: 69
    Object                  Name                    Type            Index
    Classifier              ipprec-compatibility    ip               13
```

show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601, Generation: 141
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
```

```
  Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
  Last flapped   : 2010-09-07 06:35:22 PDT (15:14:42 ago)
  Statistics last cleared: Never  Exclude Overhead Bytes: Disabled
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 total statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Drop   bytes  :                   0                   0 bps
   Drop   packets:                   0                   0 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 af3                            0                   0                   0
    1 af2                            0                   0                   0
    2 ef2                            0                   0                   0
    3 ef1                            0                   0                   0
  Egress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 af3                            0                   0                   0
    1 af2                            0                   0                   0
    2 ef2                            0                   0                   0
    3 ef1                            0                   0                   0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
```

```
    Total packets                          0                0
    Unicast packets                        0                0
    Broadcast packets                      0                0
    Multicast packets                      0                0
    CRC/Align errors                       0                0
    FIFO errors                            0                0
    MAC control frames                     0                0
    MAC pause frames                       0                0
    Oversized frames                       0
    Jabber frames                          0
    Fragment frames                        0
    VLAN tagged frames                     0
    Code violations                        0
  Filter statistics:
    Input packet count                     0
    Input packet rejects                   0
    Input DA rejects                       0
    Input SA rejects                       0
    Output packet count                              0
    Output packet pad count                          0
    Output packet error count                        0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
    Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth              Buffer Priority   Limit
                              %            bps     %           usec
    2 ef2                    39          19500     0            120     high    none
    Direction : Input
    CoS transmit queue               Bandwidth              Buffer Priority   Limit
                              %            bps     %           usec
    0 af3                    30           3000    45              0      low    none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
```

```
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
```

```
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
```

```
    Tail-dropped packets : Not Available
    RL-dropped packets   :                      0                    0 pps
    RL-dropped bytes     :                      0                    0 bps
    RED-dropped packets  :                      0                    0 pps
    RED-dropped bytes    :                      0                    0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                      0                    0 pps
    Bytes                :                      0                    0 bps
  Transmitted:
    Packets              :                      0                     0 pps
    Bytes                :                      0                    0 bps
    Tail-dropped packets :                      0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                      0                    0 pps
    Bytes                :                      0                    0 bps
  Transmitted:
    Packets              :                      0                     0 pps
    Bytes                :                      0                    0 bps
    Tail-dropped packets :                      0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                      0                    0 pps
    Bytes                :                      0                    0 bps
  Transmitted:
    Packets              :                      0                     0 pps
    Bytes                :                      0                    0 bps
    Tail-dropped packets :                      0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                 108546                    0 pps
    Bytes                :               12754752                  376 bps
  Transmitted:
```

```
    Packets              :                108546                      0 pps
    Bytes                :              12754752                    376 bps
    Tail-dropped packets :                     0                      0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available

Physical interface: ge-0/3/0, Index: 138
Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps

Scheduler map: interface-schedular-map, Index: 58414

  Scheduler: ef2, Forwarding class: ef2, Index: 39155
    Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer Limit: none,
Priority: high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
  Input shaping rate: 10000 bps
  Input scheduler map: schedular-map

Scheduler map: schedular-map, Index: 15103

  Scheduler: af3, Forwarding class: af3, Index: 35058
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
```

```
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         40582    green
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any         18928    yellow
Drop profile: green, Type: discrete, Index: 40582
  Fill level    Drop probability
          50                   0
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: yellow, Type: discrete, Index: 18928
  Fill level    Drop probability
          50                   0
         100                 100
  Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

  Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
```

```
  Fill level    Drop probability
         100                100

  Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100

  Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
```

```
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

  Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
  Congestion-notification: Disabled
Forwarding class                      ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                  0       0          0             low
normal
  af2                                  1       1          1             low
normal
  ef2                                  2       2          2             high
normal
  ef1                                  3       3          3             high
normal
  af1                                  4       4          0             low
normal

  Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
```

```
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                   0 bps
     Output bytes  :                   0                   0 bps
     Input  packets:                   0                   0 pps
     Output packets:                   0                   0 pps
    Protocol inet, MTU: 1500, Generation: 172, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Input Filters: filter-in-ge-0/3/0.0-i,
      Policer: Input: p1-ge-0/3/0.0-inet-i
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
      Flags: Is-Primary
      Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)
Forwarding class ID  Queue  Restricted queue  Fabric priority  Policing priority   SPU priority
best-effort      0   0      0                 low              normal               low

Aggregate Forwarding-class statistics per forwarding-class
  Aggregate Forwarding-class statistics:
   Forwarding-class statistics:

Forwarding-class best-effort statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0
```

```
Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

IPv4 protocol forwarding-class statistics:
  Forwarding-class statistics:
  Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

    IPv6  protocol forwarding-class statistics:
    Forwarding-class statistics:
     Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
```

```
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0

Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet  filter-in-ge-0/3/0.0-i
                           mpls                       exp-filter
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up
                           inet  p1-ge-0/3/0.0-inet-i
                           mpls

Filter: filter-in-ge-0/3/0.0-i
Counters:
Name                                                Bytes              Packets
count-filter-in-ge-0/3/0.0-i                            0                    0

Filter: exp-filter
Counters:
Name                                                Bytes              Packets
count-exp-seven-match                                   0                    0
count-exp-zero-match                                    0                    0
Policers:
Name                                               Packets
```

```
p1-ge-0/3/0.0-inet-i                                0


  Logical interface: ge-0/3/0.0, Index: 68
    Object                  Name                  Type                    Index
    Rewrite                 exp-default           exp (mpls-any)             33

Rewrite rule: exp-default, Code point type: exp, Index: 33
  Forwarding class                    Loss priority       Code point
  af3                                 low                 000
  af3                                 high                001
  af2                                 low                 010
  af2                                 high                011
  ef2                                 low                 100
  ef2                                 high                101
  ef1                                 low                 110
  ef1                                 high                111
    Object                  Name                  Type                    Index
    Classifier              exp-default           exp                        10

Classifier: exp-default, Code point type: exp, Index: 10
  Code point         Forwarding class                   Loss priority
  000                af3                                low
  001                af3                                high
  010                af2                                low
  011                af2                                high
  100                ef2                                low
  101                ef2                                high
  110                ef1                                low
  111                ef1                                high
    Object                  Name                  Type                    Index
    Classifier              ipprec-compatibility  ip                         13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                   Loss priority
  000                af3                                low
  001                af3                                high
  010                af3                                low
  011                af3                                high
  100                af3                                low
  101                af3                                high
  110                ef1                                low
  111                ef1                                high
```

```
Forwarding class                        ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                    0       0          0             low
normal
  af2                                    1       1          1             low
normal
  ef2                                    2       2          2             high
normal
  ef1                                    3       3          3             high
normal
  af1                                    4       4          0             low
normal

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                   0 bps
     Output bytes  :                   0                   0 bps
     Input  packets:                   0                   0 pps
     Output packets:                   0                   0 pps
    Protocol inet, MTU: 1500, Generation: 174, Route table: 0
      Flags: Sendbcast-pkt-to-re

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0


Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up
```

```
                          mpls

  Logical interface: ge-0/3/0.1, Index: 69
    Object                  Name                    Type                    Index
    Classifier              ipprec-compatibility    ip                         13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                    Loss priority
  000                af3                                 low
  001                af3                                 high
  010                af3                                 low
  011                af3                                 high
  100                af3                                 low
  101                af3                                 high
  110                ef1                                 low
  111                ef1                                 high
Forwarding class                      ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                  0       0          0             low
normal
  af2                                  1       1          1             low
normal
  ef2                                  2       2          2             high
normal
  ef1                                  3       3          3             high
normal
  af1                                  4       4          0             low
normal
```

show class-of-service interface (ACX Series Routers)

```
user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/0.0, Index: 69
```

```
  Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/1.0, Index: 71

  Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 dscp-default            dscp                      31
Classifier              d1                      dscp                   11331
Classifier              ci                      ieee8021p                583

  Logical interface: ge-0/1/0.0, Index: 73
Object                  Name                    Type                    Index
Rewrite                 custom-exp              exp (mpls-any)         46413

  Logical interface: ge-0/1/0.1, Index: 74

  Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                        13

  Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 ri                      ieee8021p (outer)       35392
```

```
Classifier             ci                    ieee8021p               583

Physical interface: ge-0/1/3, Index: 149
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                 Name                  Type                  Index
Classifier             ipprec-compatibility  ip                       13

  Logical interface: ge-0/1/3.0, Index: 77
Object                 Name                  Type                  Index
Rewrite                custom-exp2           exp (mpls-any)        53581

Physical interface: ge-0/1/4, Index: 150
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                 Name                  Type                  Index
Classifier             ipprec-compatibility  ip                       13

Physical interface: ge-0/1/5, Index: 151
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                 Name                  Type                  Index
Classifier             ipprec-compatibility  ip                       13

Physical interface: ge-0/1/6, Index: 152
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                 Name                  Type                  Index
Classifier             ipprec-compatibility  ip                       13

Physical interface: ge-0/1/7, Index: 153
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                 Name                  Type                  Index
Classifier             d1                    dscp                  11331

Physical interface: ge-0/2/0, Index: 154
Maximum usable queues: 8, Queues in use: 5
```

```
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

Physical interface: ge-0/2/1, Index: 155
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: ge-0/2/1.0, Index: 78

  Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#
```

show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
  Logical interface: pp0.3221225475, Index: 3221225475
Object                  Name                    Type                    Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output      4294967312
Scheduler-map           SMAP-1_UID1002          Output              4294967327
Rewrite-Output          ieee-rewrite            ieee8021p               60432
Rewrite-Output          rule1                   ip                      50463

    Adjusting application: PPPoE IA tags
      Adjustment type: absolute
      Configured shaping rate: 11000000
      Adjustment value: 5000000
      Adjustment target: node

    Adjusting application: ucac
      Adjustment type: delta
      Configured shaping rate: 5000000
      Adjustment value: 100000
      Adjustment target: node
```

Release Information

Command introduced before Junos OS Release 7.4.

Forwarding class map information added in Junos OS Release 9.4.

Options `detail` and `comprehensive` introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show class-of-service interface-set

IN THIS SECTION

- Syntax | 2762
- Description | 2762
- Options | 2762
- Required Privilege Level | 2762
- Output Fields | 2763
- Sample Output | 2765
- Release Information | 2765

Syntax

```
show class-of-service interface-set
<interface-set-name>
```

Description

Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical class of service (CoS).

Options

none	Display CoS associations for all logical interface sets.
interface-set *interface-set-name*	(Optional) Display CoS associations for the specified interface set.

Required Privilege Level

view

Output Fields

Table 105 on page 2763 describes the output fields for the `show class-of-service interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 105: show class-of-service interface-set Output Fields

Field Name	Field Description
`Interface-set`	Name of a logical interface set composed of one or more logical interfaces for which hierarchical scheduling is enabled.
`Index`	Index number of this interface set or the internal index number of this object.
`Physical interface`	Name of a physical interface.
`Queues supported`	Number of queues you can configure on the interface.
`Queues in use`	Number of queues currently configured.
`Output traffic control profile`	Name of the output traffic control profile attached to the logical interface set.
`Output traffic control profile remaining`	(Enhanced subscriber management for MX Series routers) For dynamic subscriber management, name of the output traffic control profile for remaining traffic attached to the logical interface set.

Table 105: show class-of-service interface-set Output Fields *(Continued)*

Field Name	Field Description
`Adjusting application`	Name of the application that communicates shaping-rate adjustment information to the Junos OS class-of-service process (**cosd**) on the broadband services router (BSR). The BSR uses the information from this application to perform shaping-rate adjustments on the scheduler node that manages the interface set. The adjusting application appears as **ancp LS-0** which is the Junos OS Access Node Control Profile process (**ancpd**) that performs shaping-rate adjustments on schedule nodes. The nodes are logical interface sets configured to represent subscriber local loops. When the synchronization speed of the DSL line changes, **ancpd** communicates the local loop speed to **cosd** over the default logical system, **LS-0**, and then the BSR throttles the shaping rate on the scheduler node to the loop speed. The adjusting application can also appear as **PPPoE**, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual data rate downstream attribute. The overhead accounting value is based on the access loop encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
`Adjustment type`	Type of shaping-rate adjustment performed by the BSR on the scheduler node. The type of adjustment appears as **Adjustment type**, meaning that the configured shaping rate is adjusted by an absolute value as opposed to by a percentage of the configured rate.
`Configured shaping rate`	The maximum transmission rate on the physical interface as configured by the output traffic-control profile attached to the scheduler node.
`Adjustment value`	Value of the shaping-rate adjustment information sent by the adjusting application to **cosd**.
`Adjustment overhead-accounting mode`	Configured shaping mode: `frame` or `cell`.

Sample Output

show class-of-service interface-set

```
user@host> show class-of-service interface-set example-ifset-ge-4/0/0-7
Interface-set: example-ifset-ge-4/0/0-7, Index: 8
Physical interface: ge-4/0/0, Index: 270
Queues supported: 8, Queues in use: 8
  Output traffic control profile: example-tcp-basic-rate, Index: 11395
Adjusting application: ancp LS-0
  Adjustment type: absolute
  Configured shaping rate: 50000000
  Adjustment value: 888000
  Adjustment overhead-accounting mode: cell
```

show class-of-service interface-set (Enhanced Subscriber Management)

```
user@host> show class of service interface-set
Interface-set: ge-1/0/0-201-201, Index: 1
Physical interface: ge-1/0/0, Index: 142
Queues supported: 8, Queues in use: 4
  Output traffic control profile: LEVEL_2_UID1001, Index: 4294967307
  Output traffic control profile remaining: TCP_REMAIN_UID1003, Index: 4294967308
```

Release Information

Command introduced in Junos OS Release 9.4.

show class-of-service scheduler-map

IN THIS SECTION

- Syntax | 2766
- Description | 2766

- Options | 2766
- Required Privilege Level | 2766
- Output Fields | 2766
- Sample Output | 2769
- Release Information | 2770

Syntax

```
show class-of-service scheduler-map
<name>
```

Description

Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options

none Display all scheduler maps.

name (Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level

view

Output Fields

Table 106 on page 2767 describes the output fields for the `show class-of-service scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 106: show class-of-service scheduler-map Output Fields

Field Name	Field Description
`Scheduler map`	Name of the scheduler map. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Index`	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
`Scheduler`	Name of the scheduler.
`Forwarding class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
`Transmit rate`	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword `remainder`, which indicates that the scheduler receives the remaining bandwidth of the interface.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are `none`, meaning no rate limiting, and `exact`, meaning the queue only transmits at the configured rate.
`Maximum buffer delay`	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword `remainder` to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.

Table 106: show class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
`Buffer size`	The size of the buffer as a percent of the total buffer size for the port. The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent. On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.
`Buffer dynamic threshold`	The alpha value of the shared-buffer pool at the egress buffer partition for each queue.
`Priority`	Scheduling priority: `low` or `high`.
`Excess priority`	Priority of excess bandwidth: `low`, `medium-low`, `medium-high`, `high`, or `none`.
`Explicit Congestion Notification`	(QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state: • Disable—ECN is disabled on the specified scheduler • Enable—ECN is enabled on the specified scheduler ECN is disabled by default.
`Adjust minimum`	Minimum shaping rate for an adjusted queue, in bps.
`Adjust percent`	Bandwidth adjustment applied to a queue, in percent.
`Drop profiles`	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
`Loss priority`	Packet loss priority for drop profile assignment.
`Protocol`	Transport protocol for drop profile assignment.

Table 106: show class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map

```
user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

  Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
   Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile

  Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
   Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile
```

show class-of-service scheduler-map (QFX Series)

```
user@switch# show class-of-service scheduler-map
Scheduler map: be-map, Index: 12240
```

```
Scheduler:be-sched, Forwarding class: best-effort, Index: 115
  Transmit rate: 30 percent, Rate Limit: none, Buffer size: remainder,
  Buffer Limit: none, Priority: low
  Excess Priority: unspecified, Explicit Congestion Notification: disable

  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any          3312    lan-dp
    Medium-high     any          2714    be-dp1
    High            any          3178    be-dp2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show class-of-service traffic-control-profile

IN THIS SECTION

- Syntax | **2771**
- Description | **2771**
- Options | **2771**
- Required Privilege Level | **2771**
- Output Fields | **2771**
- Sample Output | **2774**
- Release Information | **2776**

Syntax

```
show class-of-service traffic-control-profile
<profile-name>
```

Description

Display information on traffic shaping and scheduling profiles.

Options

none	Display all profiles.
profile-name	(Optional) Display information about a single profile.

Required Privilege Level

view

Output Fields

Table 107 on page 2771 describes the possible output fields for the `show class-of-service traffic-control-profile` command. The output fields that display depend on the platform and software device of your device. Output fields are listed in the approximate order in which they appear.

Table 107: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
`Traffic control profile`	Name of the traffic control profile. (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, `TC_PROF_100_199_SERIES_UID1000`) instead of with a subscriber interface.

Table 107: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
Index	Index number of the traffic control profile. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
ATM Service	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: • cbr—Constant bit rate. • rtvbr—Real time variable bit rate. • nrtvbr—Non real time variable bit rate. • ubr—Unspecified bit rate.
Maximum Burst Size	Configured maximum burst size, in cells.
Peak rate	Configured peak rate, in cps.
Sustained rate	Configured sustained rate, in cps.
Shaping rate	Configured shaping rate, in bps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
Shaping rate burst	Configured burst size for the shaping rate, in bytes. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
Shaping rate priority high	Configured shaping rate for high-priority traffic, in bps.
Shaping rate priority medium	Configured shaping rate for medium-priority traffic, in bps.

Table 107: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
Shaping rate priority low	Configured shaping rate for low-priority traffic, in bps.
Shaping rate excess high	Configured shaping rate for high-priority excess traffic, in bps.
Shaping rate excess low	Configured shaping rate for low-priority excess traffic, in bps.
Scheduler map	Name of the associated scheduler map. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.
Delay Buffer rate	Configured delay buffer rate, in bps.
Excess rate	Configured excess rate, in percent or proportion.
Excess rate high	Configured excess rate for high priority traffic, in percent or proportion.
Excess rate low	Configured excess rate for low priority traffic, in percent or proportion.
Guaranteed rate	Configured guaranteed rate, in bps or cps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: • cbr—Guaranteed rate is equal to the configured peak rate in cps. • rtvbr—Guaranteed rate is equal to the configured sustained rate in cps. • nrtvbr—Guaranteed rate is equal to the configured sustained rate in cps.
Guaranteed rate burst	Configured burst size for the guaranteed rate, in bytes.
adjust-minimum	Configured minimum shaping rate for an adjusted queue, in bps.

Table 107: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`overhead accounting mode`	Configured shaping mode: `Frame Mode` or `Cell Mode`.
`Overhead bytes`	Configured byte adjustment value.
`Adjust parent`	Configured shaping-rate adjustment for parent scheduler nodes. If enabled, this field appears. `flow-aware` indicates that the parent scheduler node is adjusted only once per multicast channel.

Sample Output

show class-of-service traffic-control-profile

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000
  Overhead bytes: 20

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
 .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
```

```
  Guaranteed rate: 4000000
..adjust-minimum 20000000

 Traffic control profile: foo, Index: 57626
   Shaping rate: 100000000
   Scheduler map: <default>
   Overhead accounting mode: Frame Mode
   Frame mode overhead accounting bytes: -12
   Adjust parent: flow-aware
```

show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
  Shaping rate burst: 500 cells
  Delay Buffer rate: 2000 cps
  Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
  ATM Service: UBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
```

show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
  ATM Service: RTVBR
  Shaping rate: 2000 cps
  Shaping rate burst: 200 cells
  Scheduler map: <default>
  Delay Buffer rate: 1000 cps
  Guaranteed rate: 1700 cps
```

show class-of-service traffic-control-profile (Enhanced Subscriber Management)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: TC_PROF_100_199_SERIES_UID1000, Index: 4294967313
  Shaping rate: 11000000
  Shaping rate burst: 1 bytes
  Scheduler map: SMAP-1_UID1002
  Delay Buffer rate: 5000000
  Overhead accounting mode: Cell Mode
  Frame mode overhead accounting bytes: -4
  Cell mode overhead accounting bytes: 20
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show database-replication statistics

IN THIS SECTION

- Syntax | 2777
- Description | 2777
- Options | 2777
- Required Privilege Level | 2777
- Output Fields | 2777
- Sample Output | 2778
- Release Information | 2778

Syntax

```
show database-replication statistics
```

Description

Display statistics regarding the replication of the subscriber management session database.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 108 on page 2777 lists the output fields for the `show database-replication statistics` command. Output fields are listed in the approximate order in which they appear.

Table 108: show database-replication statistics Output Fields

Field Name	Field Description
General	Number of dropped connections and the maximum buffer count.
Message Received	Total size of messages received and the number of received messages that have been processed.
Message Sent	Total size of messages sent and the number of sent messages that have been processed.
Message Queue	Number of messages in the queue and the maximum size of the queue.

Sample Output

show database-replication statistics

```
user@host> show database-replication statistics

General:
    Dropped connections        0
    Max buffer count           0
Message received:
    Size (bytes)               0
    Processed                  0
Message sent:
    Size (bytes)               0
    Processed                  0
Message queue:
    Queue full                 0
    Max queue size             0
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show database-replication summary | 2778

show database-replication summary

IN THIS SECTION

- Syntax | 2779
- Description | 2779
- Options | 2779

- Required Privilege Level | 2779
- Output Fields | 2779
- Sample Output | 2781
- Release Information | 2782

Syntax

```
show database-replication summary
```

Description

Display summary information regarding database replication for the subscriber management session database.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 109 on page 2780 lists the output fields for the `show database-replication summary` command. Output fields are listed in the approximate order in which they appear.

Table 109: show database-replication summary Output Fields

Field Name	Field Description
Graceful Restart	State of graceful Routing Engine switchover (GRES): • Enabled • Disabled
Mastership	State of the Routing Engine primary role: • Master • Standby
Connection	State of the connection: • Up • Down
Disconnection Reason	Reason the Routing Engines are disconnected. • RE DRAM Size Mismatch—Displayed when the amount of DRAM is different on the primary and standby Routing Engines.
Database	State of the subscriber management database: • Available • Unavailable • Synchronized

Table 109: show database-replication summary Output Fields *(Continued)*

Field Name	Field Description
Message Queue	State of the message queue: • Full • Init • Not Ready • Ready

Sample Output

show database-replication summary

```
user@host> show database-replication summary
General:
    Graceful Restart     Enabled
    Mastership           Standby
    Connection           Up
    Database             Available
    Message Queue        Ready
```

show database-replication summary (DRAM Size Mismatch Error)

```
user@host> show database-replication summary
General:
    Graceful Restart     Enabled
    Mastership           Standby
    Connection                Down
    Disconnection Reason      RE DRAM Size Mismatch
    Database                  Available
    Message Queue             Not Ready
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show database-replication statistics | 2776

show dhcp relay active-leasequery

IN THIS SECTION

- Syntax | 2782
- Description | 2782
- Options | 2783
- Required Privilege Level | 2783
- Output Fields | 2783
- Sample Output | 2786
- Release Information | 2788

Syntax

```
show dhcp relay active-leasequery
<details | summary>
<interface interface-name>
<logical-system logical-system-name>
<peer ip-address>
<routing-instance routing-instance-name>
<statistics>
```

Description

Display information about DHCPv4 active leasequery peer relay agents.

Options

details (Optional) Display the topology discovery translation table for the specified peer. You must also specify the `peer` *`ip-address`* option.

summary (Optional) Display summary information for all active leasequery peers. The `summary` option produces the same output as not specifying any option. You can also specify the `logical-system` *`logical-system-name`* option or the `routing-instance` *`routing-instance-name`* option with the `summary` option. You cannot specify any other option with the `summary` option.

interface *interface-name* (Optional) Display active leasequery statistics for a specific access interface. You must also specify the `statistics` option.

logical-system *logical-system-name* (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

peer *ip-address* (Optional) Display information about active leasequery peer relay agents. You must also specify either the `details` option or the `statistics` option.

routing-instance *routing-instance-name* (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

statistics (Optional) Display active leasequery statistics for a specific active leasequery peer relay agent or a specific access interface. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

Required Privilege Level

view

Output Fields

Table 110 on page 2784 lists the output fields for the `show dhcp relay active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 110: show dhcp relay active-leasequery Output Fields

Field Name	Field Description
peer	IP address of topology discovery peer relay agent.
Connected Peers	Number of access interfaces that have completed the topology discovery process and are associated with the peer. Each interface is in the Connected state.
Connecting Peers	Number of access interfaces that are configured with topology discovery but have not completed the topology discovery process. Consequently these interfaces are not yet associated with any peer. Each interface is in the Connecting state. When topology discovery completes and the interface is associated with a peer, the state moves to Connected.
Local Circuit ID	The local access interface for the specified peer.
Remote Circuit ID	The remote access interface on another peer relay agent that corresponds to the specified peer's local access interface.
Local Interface Address	IP address of the local access interface.
State	State of the topology discovery process. • Done—Topology discovery has completed for the specified peer.
Redundancy State	M:N redundancy state associated with the peer or interface. • Backup—The specified peer or interface is currently in backup mode. This means that the BNG hosting the relay agent is the current backup BNG for the group. • Master—The specified peer or interface is currently in primary mode. This means that the BNG hosting the relay agent is the current primary BNG for the group. • Unknown—The redundancy state of the peer or interface is unknown.
xid	Randomly generated, temporary transaction ID for the topology discovery query sent for the local access interface.

Table 110: show dhcp relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
`Remote ALQ Status`	State of the active leasequery process. • `Done`—The active leasequery process has completed. Subscriber state and binding information has been synchronized for subscriber groups that use the local access interface. • `Queued`—The active leasequery is queued. Subscriber state and binding information are not yet synchronized/ • `Unknown`—Active leasequery process state is unknown.
`Interface`	Name of the specified address for which statistics are displayed.
`Topology-Discover Configured`	Indicates whether topology discovery has been configured on the specified peer or interface, `Yes` or `No`.
`Bindings Sent`	Number of DHCP bindings sent based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings sent over the interface. For the peer, it's the count of all bindings sent over all interfaces that belong to the peer.
`Bindings Received`	Number of DHCP bindings received based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings received over the interface. For the peer, it's the count of all bindings received over all interfaces that belong to the peer.
`Bindings Installed Successfully`	Number of DHCP bindings successfully installed based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings successfully installed over the interface. For the peer, it's the count of all bindings successfully installed over all interfaces that belong to the peer.

Table 110: show dhcp relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
Bindings Failed to Install	Number of DHCP bindings that failed to install based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings that failed over the interface. For the peer, it's the count of all bindings that failed over all interfaces that belong to the peer.
Last Synchronization Time	Time stamp for the last time synchronization for subscribers occurred for the specified interface or peer.
ALQ Transmit Buffer count	The maximum TCP received and transmitted bytes buffer for the jdhcpd process. The value is nonconfigurable and is always 65,535 (ffff).
Max Leasequery Transmit Rate	The transmit rate of bulk leasequery replies per second. It is the maximum number of bulk leasequery replies per one-second interval. The value is nonconfigurable and is always 60. For example, if 600 leases are synchronized to the peer, it takes 10 seconds to complete.
Local Interface Count	Number of local access interfaces for the specified peer or interface. For the interface, the value is always 1.
Remote Interface Count	Number of remote access interfaces for the specified peer or interface. For the interface, the value is always 1.

Sample Output

show dhcp relay active-leasequery (Summary)

```
user@host> show dhcp relay active-leasequery summary

            Active Lease-query peer List:

            peer                 Connected Peers         Connecting Peers
```

```
            192.0.2.2             1000                    0
            198.51.100.2          1000                    0
            10:80:3::2            2                       0
```

show dhcp relay active-leasequery (IPv4 Peer Details)

```
user@host> show dhcp relay active-leasequery peer 192.0.2.2 details

 Local Circuit-ID   Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid        Remote ALQ Status
 ge-2/1/1.1020      ge-0/0/6.1020       192.0.2.235               Done    Backup
0xb9dc35   Done
```

show dhcp relay active-leasequery (IPv6 Peer Details)

```
user@host> show dhcp relay active-leasequery peer 2001:db8:1:3::20 details

Local Circuit-ID    Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid         Remote ALQ Status
 ps3.0               ps3.0              2001:db8:80:1::1          Done    Master
0x663345    Done
 ps4.0               ps4.0              2001:db8:90:1::2          Done    Master
0xb68623    Done
```

show dhcp relay active-leasequery (Peer Details Pseudowire Interfaces)

```
user@host> show dhcp relay active-leasequery statistics interface peer 198.51.100.1 details

 Local Circuit-ID    Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid         Remote ALQ Status
 ps1.0               ps1.0               198.51.100.1              Done    Backup
0x5da771    Done
 ps2.0               ps2.0               198.51.100.1              Done    Backup
0x7c7859    Unknown
```

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2 routing-instance RI-test-vrf

     peer : 192.0.2.2
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 111
     Bindings Received                     : 0
     Bindings Installed Successfully       : 0
     Bindings Failed to install            : 0
     Last Synchronization Time             : 2019-02-15 16:28:36 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
     Local Interface count                 : 4
     Remote Interface count                : 4
```

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020 routing-instance R1-test-vrf

     Interface : ge-2/1/1.1020
     Topology-Discover Configured          : Yes
     State                                 : Done
     Bindings Sent                         : 400
     Bindings Received                     : 0
     Bindings Installed Successfully       : 0
     Bindings Failed to install            : 0
     Last Synchronization Time             : 2019-02-15 16:20:05 IST
     ALQ Transmit Buffer count             : ffff
     Max Leasequery Transmit Rate          : 60
     Local Interface count                 : 1
     Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear dhcp relay active-leasequery statistics | 2549

show dhcpv6 relay active-leasequery | 2831

show dhcp relay binding

IN THIS SECTION

- Syntax | 2789
- Description | 2789
- Options | 2790
- Required Privilege Level | 2790
- Output Fields | 2790
- Sample Output | 2793
- Release Information | 2796

Syntax

```
show dhcp relay binding
<address>
<brief>
<detail>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<ip-address | mac-address>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address	(Optional) Display DHCP binding information for a specific client identified by one of the following entries: • *ip-address*—The specified IP address. • *mac-address*—The specified MAC address. • *session-id*—The specified session ID.
brief	(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcp relay binding`.
detail	(Optional) Display detailed client binding information.
interface *interface-name*	(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.
interfaces-vlan	(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance.
summary	(Optional) Display a summary of DHCP client information.

Required Privilege Level

view

Output Fields

Table 111 on page 2791 lists the output fields for the `show dhcp relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 111: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail
State	State of the DHCP relay address binding table on the DHCP client: • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers.	brief detail

Table 111: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface`	Incoming client interface.	brief
`Lease Expires`	Date and time at which the client's IP address lease expires.	detail
`Lease Expires in`	Number of seconds in which the lease expires.	detail
`Lease Start`	Date and time at which the client's IP address lease started.	detail
`Lease time violated`	Lease time violation has occurred.	detail
`Incoming Client Interface`	Client's incoming interface.	detail
`Server IP Address`	IP address of the DHCP server.	detail
`Server Interface`	Interface of the DHCP server.	detail
`Bootp Relay Address`	IP address of BOOTP relay.	detail
`Type`	Type of DHCP packet processing performed on the router: • `active`—Router actively processes and relays DHCP packets. • `passive`—Router passively snoops DHCP packets passing through the router.	All levels
`Lease expires at`	Date and time at which the client's IP address lease expires.	All levels

Table 111: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Dual Stack Group	Name of dual stack that is configured with the DHCP binding.	detail
Dual Stack Peer Prefix	Prefix of dual stack DHCPv6 peer.	detail
Dual Stack Peer Address	Address of the dual stack DHCPv6 peer.	detail

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.11     41          00:00:5e:00:53:01  86371       BOUND      ge-1/0/0.0
198.51.100.12     42          00:00:5e:00:53:02  86371       BOUND      ge-1/0/0.0
198.51.100.13     43          00:00:5e:00:53:03  86371       BOUND      ge-1/0/0.0
198.51.100.14     44          00:00:5e:00:53:04  86371       BOUND      ge-1/0/0.0
198.51.100.15     45          00:00:5e:00:53:05  86371       BOUND      ge-1/0/0.0
```

show dhcp relay binding detail

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


Client IP Address:  198.51.100.11
     Hardware Address:             00:00:5e:00:53:01
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
     Lease Start:                  2009-07-20 11:00:06 PDT
     Lease time violated:          yes
     Last Packet Received:         2009-07-20 11:00:06 PDT
```

```
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.22
     Server Interface:             none
     Bootp Relay Address:          198.51.100.32
     Session Id:                   41
     Dual Stack Group:             dual-stack-retail6
     Dual Stack Peer Prefix:       2001:db8:0:4::/64
     Dual Stack Peer Address:      2001:db8:1:0:8003::1/128

Client IP Address:  198.51.100.12
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
     Lease Start:                  2009-07-20 11:00:06 PDT
     Last Packet Received:         2009-07-20 11:00:06 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.22
     Server Interface:             none
     Bootp Relay Address:          198.51.100.32
     Session Id:                   42
     Generated Remote ID           host:ge-1/0/0:100
```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2

IP address       Hardware address   Type     Lease expires at
198.51.100.1      00:00:5e:00:53:01  active   2007-03-27 15:06:20 EDT
```

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100

IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:94  86124       BOUND      ge-1/1/0:100
```

show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100

IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.16      7           00:00:5e:00:53:92  86124       BOUND      ge-1/1/0:10-100
```

show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 198.51.100.13
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.13      43          00:00:5e:00:53:03  86293       BOUND      ge-1/0/0.0
```

show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:00:5e:00:53:05
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      45          00:00:5e:00:53:05  86279       BOUND      ge-1/0/0.0
```

show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
 IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.11      41          00:00:5e:00:53:53  86305       BOUND      ge-1/0/0.0
```

show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17       42          00:00:5e:00:53:02  86346       BOUND      ge-1/0/0.1073741827
192.168.0.16       41          00:00:5e:00:53:01  86346       BOUND      ge-1/0/0.1073741827
```

show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9        24          00:00:5e:00:53:04  86361       BOUND      ge-1/3/0.110
192.168.0.8        23          00:00:5e:00:53:03  86361       BOUND      ge-1/3/0.110
192.168.0.7        22          00:00:5e:00:53:02  86361       BOUND      ge-1/3/0.110
```

show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)
```

Release Information

Command introduced in Junos OS Release 8.3.

Options `interface` and *mac-address* added in Junos OS Release 8.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

clear dhcp relay binding | 2553

show dhcp relay lockout-entries

IN THIS SECTION

- Syntax | 2797
- Description | 2797
- Options | 2797
- Required Privilege Level | 2797

- Output Fields | **2797**
- Sample Output | **2799**
- Release Information | **2800**

Syntax

```
show dhcp relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 relay agent lockout database.

Options

all	Display all client entries in the lockout database.
index	Number identifying a client entry to be displayed.

Required Privilege Level

view

Output Fields

Table 112 on page 2798 lists the output fields for the `show dhcp relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 112: show dhcp relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
Index	Number identifying a specific entry in the lockout database.	all and index
Key	Client identifier for the client in the lockout database.	all and index
State	Type of lockout period for the entry: • Grace—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • Lockout—Client is currently locked out; attempts to establish a session are rejected.	all and index
Expires (s)	Number of seconds until the current lockout period expires.	all only
Elapsed (s)	Number of seconds since the current lockout or grace timer started.	all only
Count	Number of consecutive times the client has been locked out.	all only
Expires	Date and time when the current lockout period ends.	index only
Expires in	Number of seconds until the current period expires.	index only
Lockout count	Number of consecutive times client has been locked out.	index only
Next lockout time	Duration of the next lockout period for this client.	index only

Table 112: show dhcp relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp relay lockout-entries (All Entries)

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp relay lockout-entries (Specific Entry)

```
user@host> show dhcp relay lockout-entries index 2
     Index:                              2
     Key:                                default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                              Lockout
     Expires:                            2018-03-29 19:06:17 IST
     Expires in:                         87
     Lockout count:                      1
     Next lockout time:                  200
     Min lockout time:                   100
     Lockout reason:                     181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

clear dhcp relay lockout-entries | 2556
show dhcp server lockout-entries | 2821
show dhcpv6 relay lockout-entries | 2850
show dhcpv6 server lockout-entries | 2874

show dhcp relay statistics

IN THIS SECTION

- Syntax | 2800
- Syntax | 2801
- Description | 2801
- Options | 2801
- Required Privilege Level | 2801
- Output Fields | 2801
- Sample Output | 2805
- Release Information | 2806

Syntax

```
show dhcp relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Syntax

Syntax for EX Series switches:

```
show dhcp relay statistics
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCP relay bulk leasequery statistics.
leasequery	(Optional) Display information about DHCP relay individual leasequery statistics.
logical-system *logical-system-name*	(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 113 on page 2802 lists the output fields for the `show dhcp relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 113: show dhcp relay statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the `Packets dropped` output. When all of the Packets dropped statistics are 0 (zero), only the `Total` field appears. • `Total`—Total number of packets discarded by the extended DHCP relay agent application. • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified. • `Bad opcode`—Number of packets discarded because an invalid operation code was specified. • `Bad options`—Number of packets discarded because invalid options were specified. • `Invalid server address`—Number of packets discarded because an invalid server address was specified. • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment. • `No interface match`—Number of packets discarded because they did not belong to a configured interface. • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance. • `No valid local address`—Number of packets discarded because there was no valid local address. • `Packet too short`—Number of packets discarded because they were too short. • `Read error`—Number of packets discarded because of a system read error. • `Send error`—Number of packets that the extended DHCP relay application could not send. • `Option 60`—Number of packets discarded containing DHCP option 60 vendor-specific information. • `Option 82`—Number of packets discarded because DHCP option 82 information could not be added.

Table 113: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
Messages received	Number of DHCP messages received. • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received • DHCPLEASEACTIVE—Number of active DHCP leases • DHCPLEASEUNASSIGNED—Number of DHCP leases that are managed by the server but have not yet been assigned • DHCPLEASEUNKNOWN—Number of unknown DHCP leases • DHCPLEASEQUERYDONE—The leasequery is complete
Messages sent	Number of DHCP messages sent. • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted • DHCPLEASEQUERY—Number of DHCP leasequery messages transmitted • DHCPLEASEBULKLEASEQUERY—Number of DHCP bulk leasequery messages transmitted
External Server Response	State of the external DHCP server responsiveness.

Table 113: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
Packets forwarded	Number of packets forwarded. • BOOTREQUEST—Number of BOOTREQUEST protocol data units (PDUs) forwarded • BOOTREPLY—Number of BOOTREPLY protocol data units (PDUs) forwarded
External Server Response	State of the external DHCP server responsiveness.
Total Requested Servers	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
Total Attempted Servers	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.
Total Connected	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.
Total Terminated by Server	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
Total Max Attempted	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
Total Closed due to Errors	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
In-Flight Connected	Number of current bulk leasequery connections on the DHCP relay agent.
Bulk Leaasequery Reply Packet Retries	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                       34
    Bad hardware address        1
    Bad opcode                  1
    Bad options                 3
    Invalid server address      5
    Lease Time Violation        1
    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
    DHCPLEASEACTIVE             0
    DHCPLEASEUNASSIGNED         0
    DHCPLEASEUNKNOWN            0
    DHCPLEASEQUERYDONE          0

Messages sent:
    BOOTREPLY                   0
    DHCPOFFER                   2
    DHCPACK                     1
    DHCPNAK                     0
    DHCPFORCERENEW              0
    DHCPLEASEQUERY              0
    DHCPBULKLEASEQUERY          0
```

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

External Server Response:
    State                     Responding
```

show dhcp relay statistics bulk-leasequery-connections

```
user@host> show dhcp relay statistics bulk-leasequery-connections

Total Requested Servers:    0
Total Attempted Servers:    0
Total Connected:            0
Total Terminated by Server: 0
Total Max Attempted:        0
Total Closed due to Errors: 0
In-Flight Connected:        0
Bulk Leasequery Reply Packet Retries:   0
```

Release Information

Command introduced in Junos OS Release 8.3.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcp relay statistics | 2559

show dhcp server active-leasequery

IN THIS SECTION

- Syntax | 2807
- Description | 2807
- Required Privilege Level | 2807
- Output Fields | 2807
- Sample Output | 2808
- Release Information | 2808

Syntax

```
show dhcp server active-leasequery summary
```

Description

Display the active leasequery status summary of the DHCP local server.

Required Privilege Level

view

Output Fields

Table 114 on page 2807 lists the output fields for the `show dhcp server active-leasequery summary` command.

Table 114: show dhcp server active-leasequery summary Output Fields

Field Name	Field Description
`Connected Peers`	Number of peer DHCP servers connected.

Table 114: show dhcp server active-leasequery summary Output Fields *(Continued)*

Field Name	Field Description
Connecting Peers	Number of DHCP server peers trying to connect or reconnect after failure.
peer	IP address of the peer DHCP server.

Sample Output

show dhcp server active-leasequery summary

```
user@host> show dhcp server active-leasequery summary
```

```
Active Lease-query peer List:

peer         Connected Peers   Connecting Peers
192.0.2.0    2                 0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

M:N Subscriber Service Redundancy on DHCP Server | 1219

show dhcp server active-leasequery statistics

IN THIS SECTION

- Syntax | 2809
- Description | 2809
- Options | 2809
- Required Privilege Level | 2809
- Output Fields | 2810
- Sample Output | 2811
- Release Information | 2811

Syntax

```
show dhcp server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Display the active leasequery statistics of the DHCP local server.

Options

peer *peer-address*	IP address of the DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

view

Output Fields

Table 115 on page 2810 lists the output fields for the `show dhcp server active-leasequery statistics` command.

Table 115: show dhcp server active-leasequery statistics Output Fields

Field Name	Field Description
`ALQ Transmit Buffer count`	Numbers of packets the active leasequery transmitted.
`Bindings Failed to install`	Number of binding installation attempt failed.
`Bindings Installed Successfully`	Number of successful binding installation.
`Bindings Received`	Number of binding request received.
`Bindings Sent`	Number of binding request sent.
`Last Synchronization Time`	Last synchronization time.
`Local Interface count`	Number of local interfaces.
`Max Leasequery Transmit Rate`	Maximum leasequery transmit rate.
`peer`	IP address of the peer DHCP server.
`Remote Interface count`	Number of remote interfaces.
`State`	Status of the topology discovery configuration.
`Topology-Discover Configured`	Status of the topology discovery.

Sample Output

show dhcp server active-leasequery peer 192.0.2.0

```
user@host> show dhcp server active-leasequery peer 192.0.2.0
```

```
peer : 192.0.2.0
Topology-Discover Configured          : Yes
State                                 : Done
Bindings Sent                         : 0
Bindings Received                     : 6
Bindings Installed Successfully       : 0
Bindings Failed to install            : 0
Last Synchronization Time             : 2021-04-18 19:02:16 IST
ALQ Transmit Buffer count             : 0x ffff
Max Leasequery Transmit Rate          : 60
Local Interface count                 : 1
Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear dhcp server active-leasequery statistics | 2562

show dhcp server binding

IN THIS SECTION

Syntax | 2812

Description | 2812

- Options | 2812
- Required Privilege Level | 2813
- Output Fields | 2813
- Sample Output | 2817
- Release Information | 2821

Syntax

```
show dhcp server binding
<address>
<interfaces-vlan><brief | detail | summary>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (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.

Options

address (Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.

- *session-id*—The specified session ID.

brief | detail | summary — (Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcp server binding`.

interface *interface-name* — (Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan — (Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard — (Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name* — (Optional) Display information about active client bindings for DHCP clients on the specified logical system.

routing-instance *routing-instance-name* — (Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 116 on page 2813 lists the output fields for the `show dhcp server binding` command. Output fields are listed in the approximate order in which they appear.

Table 116: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail

Table 116: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: • BOUND—Client has active IP address lease. • FORCERENEW—Client has received forcerenew message from server. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCP server. • SELECTING—Client receiving offers from DHCP servers.	brief detail
Interface	Interface on which the request was received.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which lease expires.	detail

Table 116: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lease Start`	Date and time at which the client's IP address lease started.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Last Packet Received`	Date and time at which the router received the last packet.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Client Interface Svlan Id`	S-VLAN ID of the client's incoming interface.	`detail`
`Client Interface Vlan Id`	VLAN ID of the client's incoming interface.	`detail`
`Demux Interface`	Name of the IP demultiplexing (demux) interface.	`detail`
`Server IP Address` or `Server Identifier`	IP address of DHCP server.	`detail`
`Server Interface`	Interface of DHCP server.	`detail`
`Client Pool Name`	Name of address pool used to assign client IP address lease.	`detail`

Table 116: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Liveness Detection State`	State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session: **NOTE**: This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber. • `DOWN`—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The `DOWN` state is reported only during this transition period of time. • `UNKNOWN`—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The `UNKNOWN` state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the `UP` state. • `UP`—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have *both* determined that the liveness detection session for the BFD protocol is in the `UP` state. • `WENT_DOWN`—State is functionally equivalent to the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state implying a liveness detection failure.	`detail`

Table 116: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	The WENT_DOWN state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine.	
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	detail
ACI Interface Set Index	Index number of the dynamic ACI interface set.	detail
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	detail
Client Profile Name	DHCP client profile name.	detail
Dual Stack Group	DHCP server profile name.	detail
Dual Stack Peer Prefix	IPv6 prefix of peer.	detail
Dual Stack Peer Address	IPv6 address of peer.	detail

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15     6           00:00:5e:00:53:01  86180       BOUND      ge-1/0/0.0
198.51.100.16     7           00:00:5e:00:53:02  86180       BOUND      ge-1/0/0.0
198.51.100.17     8           00:00:5e:00:53:03  86180       BOUND      ge-1/0/0.0
```

```
198.51.100.18     9           00:00:5e:00:53:04  86180       BOUND      ge-1/0/0.0
198.51.100.19     10          00:00:5e:00:53:05  86180       BOUND      ge-1/0/0.0
```

show dhcp server binding detail

```
user@host> show dhcp server binding detail
Client IP Address:  198.51.100.15
     Hardware Address:             00:00:5e:00:53:01
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.9
     Server Interface:             none
     Session Id:                   6
     Client Pool Name:             6
     Liveness Detection State:     UP
 Client IP Address:            198.51.100.16
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Lease time violated:          yes
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.9
     Server Interface:             none
     Session Id:                   7
     Client Pool Name:             7
     Liveness Detection State:     UP
```

When DHCP binding is configured with dual-stack, we get the following output:

```
user@host> show dhcp server binding detail
Client IP Address:  198.51.100.10
     Hardware Address:             00:00:64:03:01:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
     Protocol-Used:                DHCP
     Lease Expires:                2016-11-07 08:30:39 PST
```

```
    Lease Expires in:             43706 seconds
    Lease Start:                  2016-11-04 11:00:37 PDT
    Last Packet Received:         2016-11-06 09:00:39 PST
    Incoming Client Interface:    ae0.3221225472
    Client Interface Svlan Id:    2000
    Client Interface Vlan Id:     1
    Server Ip Address:            198.51.100.2
    Session Id:                   2
    Client Pool Name:             my-v4-pool
    Client Profile Name:          dhcp-retail
    Dual Stack Group:             my-dual-stack
    Dual Stack Peer Prefix:       2001:db8:ffff:0:4::/64
    Dual Stack Peer Address:      2001:db8:0:8003::1/128
```

show dhcp server binding detail (ACI Interface Set Configured)

```
user@host> show dhcp server binding detail
Client IP Address:  198.51.100.14
    Hardware Address:             00:00:5e:00:53:02
    State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
    Lease Expires:                2012-03-13 09:53:32 PDT
    Lease Expires in:             82660 seconds
    Lease Start:                  2012-03-12 10:23:32 PDT
    Last Packet Received:         2012-03-12 10:23:32 PDT
    Incoming Client Interface:    demux0.1073741827
    Client Interface Svlan Id:    1802
    Client Interface Vlan Id:     302
    Demux Interface:              demux0.1073741832
    Server Identifier:            198.51.100.202
    Session Id:                   11
    Client Pool Name:             poolA
    Client Profile Name:          DEMUXprofile
    Liveness Detection State:     UP
    ACI Interface Set Name:       aci-1002-demux0.1073741827
    ACI Interface Set Index:      2
    ACI Interface Set Session ID: 6
```

show dhcp server binding interface <vlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:100
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/1/0:100
```

show dhcp server binding interface <svlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.16      7          00:00:5e:00:53:02  86124       BOUND      ge-1/1/0:10-100
```

show dhcp server binding <ip-address>

```
user@host> show dhcp server binding 198100.19
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.19      10          00:00:5e:00:53:05  86081       BOUND      ge-1/0/0.0
```

show dhcp server binding <session-id>

```
user@host> show dhcp server binding 6
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/0/0.0
```

show dhcp server binding summary

```
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcp server binding <interfaces-vlan>

```
user@host> show dhcp server binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
```

```
192.168.0.17      42          00:00:5e:00:53:02  86346      BOUND      ge-1/0/0.1073741827
192.168.0.16      41          00:00:5e:00:53:01  86346      BOUND      ge-1/0/0.1073741827
```

show dhcp server binding <interfaces-wildcard>

```
user@host> show dhcp server binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9       24          00:00:5e:00:53:04  86361       BOUND      ge-1/3/0.110
192.168.0.8       23          00:00:5e:00:53:03  86361       BOUND      ge-1/3/0.110
192.168.0.7       22          00:00:5e:00:53:02  86361       BOUND      ge-1/3/0.110
```

Release Information

Command introduced in Junos OS Release 9.0.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

DHCP Monitoring and Management | 895

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

clear dhcp server binding | 2563

show dhcp server lockout-entries

IN THIS SECTION

- Syntax | 2822
- Description | 2822
- Options | 2822
- Required Privilege Level | 2822
- Output Fields | 2822
- Sample Output | 2824

- Release Information | 2824

Syntax

```
show dhcp server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 local lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 117 on page 2822 lists the output fields for the `show dhcp server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 117: show dhcp server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`

Table 117: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Key	Client identifier for the client in the lockout database.	all and index
State	Type of lockout period for the entry: • Grace—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • Lockout—Client is currently locked out; attempts to establish a session are rejected.	all and index
Expires (s)	Number of seconds until the current lockout period expires.	all only
Elapsed (s)	Number of seconds since the current lockout or grace timer started.	all only
Count	Number of consecutive times the client has been locked out.	all only
Expires	Date and time when the current lockout period ends.	index only
Expires in	Number of seconds until the current period expires.	index only
Lockout count	Number of consecutive times client has been locked out.	index only
Next lockout time	Duration of the next lockout period for this client.	index only
Min lockout time	Minimum duration for a lockout period; the initial lockout time.	index only

Table 117: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp server lockout-entries (All Entries)

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp server lockout-entries (Specific Entry)

```
user@host> show dhcp server lockout-entries index 2
     Index:                              2
     Key:                                default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                              Lockout
     Expires:                            2018-03-29 19:06:17 IST
     Expires in:                         87
     Lockout count:                      1
     Next lockout time:                  200
     Min lockout time:                   100
     Lockout reason:                     181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

clear dhcp server lockout-entries | 2568

show dhcp relay lockout-entries | 2796

show dhcpv6 relay lockout-entries | 2850

show dhcpv6 server lockout-entries | 2874

show dhcp server statistics

IN THIS SECTION

- Syntax | 2825
- Description | 2825
- Options | 2826
- Required Privilege Level | 2826
- Output Fields | 2826
- Sample Output | 2830
- Release Information | 2831

Syntax

```
show dhcp server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCP local server bulk leasequery statistics.
logical-system *logical-system-name*	(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 118 on page 2827 lists the output fields for the `show dhcp server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 118: show dhcp server statistics Output Fields

Field Name	Field Description
Packets dropped	Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • Total—Total number of packets discarded by the extended DHCP local server • Authentication—Number of packets discarded because they could not be authenticated • Bad hardware address—Number of packets discarded because an invalid hardware address was specified • Bad opcode—Number of packets discarded because an invalid operation code was specified • Bad options—Number of packets discarded because invalid options were specified • Dynamic profile—Number of packets discarded due to dynamic profile information • Invalid server address—Number of packets discarded because an invalid server address was specified • Lease Time Violation—Number of packets discarded because of a lease time violation • No available addresses—Number of packets discarded because there were no addresses available for assignment • No interface match—Number of packets discarded because they did not belong to a configured interface • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance • No valid local address—Number of packets discarded because there was no valid local address • Packet too short—Number of packets discarded because they were too short • Read error—Number of packets discarded because of a system read error • Send error—Number of packets that the extended DHCP local server could not send

Table 118: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
`Offer Delay`	Number of DHCPv4 offer messages delayed. • `DELAYED`—Number of DHCPv4 offer packets that have been sent after being delayed. • `INPROGRESS`—Number of DHCPv4 offer packets that are in the delay queue. • `TOTAL`—Total number of delayed DHCPv4 offer messages; sum of `DELAYED` and `INPROGRESS`.
`Messages received`	Number of DHCP messages received. • `BOOTREQUEST`—Number of BOOTP protocol data units (PDUs) received • `DHCPDECLINE`—Number of DHCP PDUs of type DECLINE received • `DHCPDISCOVER`—Number of DHCP PDUs of type DISCOVER received • `DHCPINFORM`—Number of DHCP PDUs of type INFORM received • `DHCPRELEASE`—Number of DHCP PDUs of type RELEASE received • `DHCPREQUEST`—Number of DHCP PDUs of type REQUEST received • `DHCPLEASEQUERY`—Number of DHCP leasequery messages received. • `DHCPBULKLEASEQUERY`—Number of DHCP bulk leasequery messages received. • `DHCPRENEW`—Number of DHCP renew messages received; subset of `DHCPREQUEST` counter. • `DHCPREBIND`—Number of DHCP rebind messages received; subset of `DHCPREQUEST` counter.

Table 118: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
Messages sent	Number of DHCP messages sent. • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted • DHCPLEASEUNASSIGNED—Number of DHCP leases that are managed by the server but have not yet been assigned • DHCPLEASEUNKNOWN—Number of unknown DHCP leases • DHCPLEASEACTIVE—Number of active DHCP leases • DHCPLEASEQUERYDONE—The leasequery is complete
Total Accepted Connections	Total number of bulk leasequery connections accepted by the server.
Total Not-Accepted Connections	Total number of bulk leasequery connections not accepted by the server.
Connections Closed due to Errors	Number of bulk leasequery connections that the server closed due to an internal error.
Connections Closed due to max-empty-replies	Number of bulk leasequery connections that the server closed because the maximum number of empty replies was reached.
In-flight Connections	Number of bulk leasequery connections on the server.

Sample Output

show dhcp server statistics

```
user@host> show dhcp server statistics
Packets dropped:
    Total                     1
    Lease Time Violation      1

Offer Delay:
    DELAYED                   3
    INPROGRESS                9
    TOTAL                    12

Messages received:
    BOOTREQUEST              25
    DHCPDECLINE              0
    DHCPDISCOVER             10
    DHCPINFORM               0
    DHCPRELEASE              4
    DHCPREQUEST              10
    DHCPRENEW                4
    DHCPREBIND               2

Messages sent:
    BOOTREPLY                20
    DHCPOFFER                10
    DHCPACK                  10
    DHCPNAK                  0
    DHCPFORCERENEW           0
```

show dhcp server statistics

```
user@host> show dhcp server statistics verbose
Packets dropped:
    Total                     0

Messages received:
    BOOTREQUEST              238
    DHCPDECLINE              0
```

```
    DHCPDISCOVER                1
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 237
    DHCPRENEW                   236
    DHCPREBIND                  0

Messages sent:
    BOOTREPLY                   20
    DHCPOFFER                   10
    DHCPACK                     10
    DHCPNAK                     0
    DHCPFORCERENEW              0
```

Release Information

Command introduced in Junos OS Release 9.0.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcp server statistics | 2570

show dhcpv6 relay active-leasequery

IN THIS SECTION

- Syntax | 2832
- Description | 2832
- Options | 2832
- Required Privilege Level | 2833
- Output Fields | 2833
- Sample Output | 2836
- Release Information | 2837

Syntax

```
show dhcpv6 relay active-leasequery
<details | summary>
<interface interface-name>
<logical-system logical-system-name>
<peer ipv6-address>
<routing-instance routing-instance-name>
<statistics>
```

Description

Display information about DHCPv6 active leasequery peer relay agents.

Options

details (Optional) Display detailed information about active leasequery peers or interfaces. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

summary (Optional) Display summary information about active leasequery peers. The `summary` option produces the same output as not specifying any option. You can also specify the `logical-system` *`logical-system-name`* option or the `routing-instance` *`routing-instance-name`* option with the `summary` option. You cannot specify any other option with the `summary` option.

interface *interface-name* (Optional) Display active leasequery statistics for a specific access interface. You must also specify the `statistics` option.

logical-system *logical-system-name* (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

peer *ipv6-address* (Optional) Display information about active leasequery peer relay agents. You must also specify either the `details` option or the `statistics` option.

routing-instance *routing-instance-name* (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

statistics (Optional) Display active leasequery statistics for a specific active leasequery peer relay agent or a specific access interface. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

Required Privilege Level

view

Output Fields

Table 119 on page 2833 lists the output fields for the `show dhcpv6 relay active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 119: show dhcpv6 relay active-leasequery Output Fields

Field Name	Field Description
`peer`	IP address of topology discovery peer relay agent.
`Connected Peers`	Number of access interfaces that have completed the topology discovery process and are associated with the peer. Each interface is in the Connected state.
`Connecting Peers`	Number of access interfaces that are configured with topology discovery but have not completed the topology discovery process. Consequently these interfaces are not yet associated with any peer. Each interface is in the Connecting state. When topology discovery completes and the interface is associated with a peer, the state moves to Connected.
`Local Circuit ID`	The local access interface for the specified peer.
`Remote Circuit ID`	The remote access interface on another peer relay agent that corresponds to the specified peer's local access interface.
`Local Interface Address`	IP address of the local access interface.
`State`	State of the topology discovery process. • `Done`—Topology discovery has completed for the specified peer.

Table 119: show dhcpv6 relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
Redundancy State	VRRP redundancy state associated with the peer or interface. • Backup—The specified peer or interface is currently in backup mode. This means that the BNG hosting the relay agent is the current backup BNG for the group. • Master—The specified peer or interface is currently in primary mode. This means that the BNG hosting the relay agent is the current primary BNG for the group. • Unknown—The VRRP redundancy state of the peer or interface is unknown.
xid	Randomly generated, temporary transaction ID for the topology discovery query sent for the local access interface.
Remote ALQ Status	State of the active leasequery process. • Done—The active leasequery process has completed. Subscriber state and binding information has been synchronized for subscriber groups that use the local access interface. • Queued—The active leasequery is queued. Subscriber state and binding information are not yet synchronized/ • Unknown—Active leasequery process state is unknown.
Interface	Name of the specified address for which statistics are displayed.
Topology-Discover Configured	Indicates whether topology discovery has been configured on the specified peer or interface, Yes or No.
Bindings Sent	Number of DHCP bindings sent based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings sent over the interface. For the peer, it's the count of all bindings sent over all interfaces that belong to the peer.

Table 119: show dhcpv6 relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
Bindings Received	Number of DHCP bindings received based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings received over the interface. For the peer, it's the count of all bindings received over all interfaces that belong to the peer.
Bindings Installed Successfully	Number of DHCP bindings successfully installed based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings successfully installed over the interface. For the peer, it's the count of all bindings successfully installed over all interfaces that belong to the peer.
Bindings Failed to Install	Number of DHCP bindings that failed to install based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings that failed over the interface. For the peer, it's the count of all bindings that failed over all interfaces that belong to the peer.
Last Synchronization Time	Time stamp for the last time synchronization for subscribers occurred for the specified interface or peer.
ALQ Transmit Buffer count	The maximum TCP received and transmitted bytes buffer for the jdhcpd process. The value is nonconfigurable and is always 65,535 (ffff).
Max Leasequery Transmit Rate	The transmit rate of bulk leasequery replies per second. It is the maximum number of bulk leasequery replies per one-second interval. The value is nonconfigurable and is always 60. For example, if 600 leases are synchronized to the peer, it takes 10 seconds to complete.
Local Interface Count	Number of local access interfaces for the specified peer or interface. For the interface, the value is always 1.

Table 119: show dhcpv6 relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
Remote Interface Count	Number of remote access interfaces for the specified peer or interface. For the interface, the value is always 1.

Sample Output

show dhcpv6 relay active-leasequery (Summary)

```
user@host> show dhcpv6 relay active-leasequery summary

            Active Lease-query peer List:

            peer                 Connected Peers          Connecting Peers
            2001:db8::fa:2       1                        0
            2001:db8::fe:3       1                        0
```

show dhcpv6 relay active-leasequery (Peer Details)

```
user@host> show dhcpv6 relay active-leasequery peer 2001:db8::fa:2 details

 Local Circuit-ID     Remote Circuit-ID    Local Interface Address   State    Redundancy State
xid          Remote ALQ Status
 ge-2/0/3.50          ge-1/3/1.35          2001:db8::8:10            Done     Backup
0x521bc8    Done
```

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2 routing-instance
RI-test-vrf

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
```

```
    Bindings Sent                          : 420
    Bindings Received                      : 0
    Bindings Installed Successfully        : 0
    Bindings Failed to install             : 0
    Last Synchronization Time              : 2019-03-22 09:15:15 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 4
    Remote Interface count                 : 4
```

show dhcpv6 relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1 routing-instance
R1-test-vrf

    Interface : ge-0/0/0.1
    Topology-Discover Configured           : Yes
    State                                  : Done
    Bindings Sent                          : 400
    Bindings Received                      : 0
    Bindings Installed Successfully        : 0
    Bindings Failed to install             : 0
    Last Synchronization Time              : 2019-02-28 14:11:32 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 1
    Remote Interface count                 : 1
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear dhcpv6 relay active-leasequery statistics | 2573

show dhcp relay active-leasequery | 2782

show dhcpv6 relay binding

IN THIS SECTION

- Syntax | 2838
- Description | 2838
- Options | 2838
- Required Privilege Level | 2839
- Output Fields | 2839
- Sample Output | 2843
- Release Information | 2850

Syntax

```
show dhcpv6 relay binding
<address>
<brief>
<detail>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

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

Options

address	(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

brief (Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcpv6 relay binding`.

detail (Optional) Display detailed client binding information.

interface *interface-name* (Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.

interfaces-vlan (Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.

interfaces-wildcard (Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name* (Optional) Perform this operation on the specified logical system.

routing-instance *routing-instance-name* (Optional) Perform this operation on the specified routing instance.

summary (Optional) Display a summary of DHCPv6 client information.

Required Privilege Level

view

Output Fields

Table 120 on page 2840 lists the output fields for the `show dhcpv6 relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 120: show dhcpv6 relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Client IPv6 Prefix	Prefix of the DHCPv6 client.	brief detail
Client IPv6 Excluded Prefix	IPv6 Prefix of the DHCP client excluded.	detail
Client DUID	DHCP for IPv6 Unique Identifier (DUID) of the client.	brief detail
Client IPv6 Address	IPv6 address assigned to the subscriber.	detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail

Table 120: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the DHCPv6 relay address binding table on the DHCPv6 client: • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RECONFIGURE—Client is broadcasting a request to reconfigure the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCPv6 server. • SELECTING—Client is receiving offers from DHCPv6 servers.	brief detail
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which the client's IPv6 prefix expires.	detail

Table 120: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lease Start`	Date and time at which the client's IP address lease started.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Server Address`	IP address of the DHCPv6 server. Displays `unknown` for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the `Next Hop Server Facing Relay` field.	`detail`
`Next Hop Server Facing Relay`	Next-hop address in the direction of the DHCPv6 server.	`detail`
`Server Interface`	Interface of the DHCPv6 server.	`detail`
`Relay Address`	IP address of the relay.	`detail`
`Client Pool Name`	Address pool that granted the client lease.	`detail`
`Client ID Length`	Length of client ID.	All levels
`Client Id`	Client ID.	All levels
`Generated Circuit ID`	Circuit ID generated by the DHCPv6 Interface-ID option (option 18)	`detail`
`Generated Remote ID Enterprise Number`	The Juniper Networks IANA private enterprise number	`detail`

Table 120: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Generated Remote ID	Remote ID generated by the DHCPv6 Remote-ID option (option 37)	detail
Dual Stack Group	Name of the dual-stack group for the DHCPv6 binding.	detail
Dual Stack Peer Address	Address of the dual-stack DHCPv4 peer.	detail

Sample Output

show dhcpv6 relay binding

```
user@host> show dhcpv6 relay binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64    1          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64    2          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6          83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

show dhcpv6 relay binding (Address)

```
user@host> show dhcp6 relay binding 2001:db8:1111:2222::/64 detail
Session Id:  1
     Client IPv6 Prefix:                  2001:db8:3c4d:15::/64
     Client DUID:                         LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
```

```
     State:                                BOUND(RELAY_STATE_BOUND)
     Lease Expires:                        2011-05-25 07:12:09 PDT
     Lease Expires in:                     77115 seconds
     Preferred Lease Expires:              2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:           600 seconds
     Lease Start:                          2011-05-24 07:12:09 PDT
     Incoming Client Interface:            ge-1/0/0.0
     Server Address:                       2001:db8:aaaa:bbbb::1
     Server Interface:                     none
     Relay Address:                        2001:db8:1111:2222::
     Client Pool Name:                     pool-25
     Client Id Length:                     14
     Client Id:                            /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail (Client ID)

```
user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001 detail
Session Id:  1
     Client IPv6 Prefix:                   2001:db8:3c4d:15::/64
     Client DUID:                          LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                BOUND(RELAY_STATE_BOUND)
     Lease Expires:                        2011-05-25 07:12:09 PDT
     Lease Expires in:                     77115 seconds
     Preferred Lease Expires:              2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:           600 seconds
     Lease Start:                          2011-05-24 07:12:09 PDT
     Lease time violated:                  yes
     Incoming Client Interface:            ge-1/0/0.0
     Server Address:                       2001:db8:aaaa:bbbb::1
     Server Interface:                     none
     Relay Address:                        2001:db8:1111:2222::
     Client Pool Name:                     pool-25
     Client Id Length:                     14
     Client Id:                            /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail

```
user@host> show dhcpv6 relay binding detail
Session Id:  1
```

```
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Lease time violated:                   yes
     Incoming Client Interface:             ge-1/0/0.0
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
     Generated Remote ID Enterprise Number: 1411
     Generated Remote ID:                   host:ge-1/0/0:100
```

show dhcpv6 relay binding detail (Dual-Stack)

```
user@host> show dhcpv6 relay binding detail
Session Id:  2
     Client IPv6 Prefix:                    2001:db8:ffff:0:4::/64
     Client IPv6 Address:                   2001:db8:3000:8003::1/128
     Client DUID:                           LL0x1-00:00:64:01:01:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                         2016-10-17 07:39:25 PDT
     Lease Expires in:                      3450 seconds
     Lease Start:                           2016-10-17 06:39:25 PDT
     Last Packet Received:                  2016-10-17 06:39:25 PDT
     Incoming Client Interface:             ae0.3221225472
     Client Interface Svlan Id:             2000
     Client Interface Vlan Id:              1
     Server Ip Address:                     2001:db8:3000::2
     Server Interface:                      none
     Client Profile Name:                   my-dual-stack
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006401/0x0102
```

```
     Dual Stack Group:                      group1
     Dual Stack Peer Address:               192.0.2.4
```

show dhcpv6 relay binding detail (Multi-Relay Topology)

```
user@host > show dhcpv6 relay binding detail
Session Id:  13
     Client IPv6 Prefix:                    2001:db8:3000:0:8001::5/128
     Client DUID:                           LL0x1-00:00:65:03:01:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                         2011-11-21 06:14:50 PST
     Lease Expires in:                      293 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-11-21 06:09:50 PST
     Incoming Client Interface:             ge-1/0/0.0
     Server Address:                        unknown
     Next Hop Server Facing Relay:          2001:db8:4000::2
     Server Interface:                      none
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006503/0x0102
```

show dhcpv6 relay binding (Session ID)

```
user@host> show dhcpv6 relay binding 41
Prefix                  Session Id   Expires   State   Interface    Client DUID
2001:db8:3c4d:15::/64   41           78837     BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
```

show dhcpv6 relay binding (Subscriber with Multiple Addresses)

```
user@host> show dhcpv6 relay binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1001::1:24/128          23          593      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:1001::1:1c/128          23          393      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:1001::1:14/128          23          193      BOUND    ge-9/0/9.0
```

```
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::300/120           23          293       BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::200/120           23          193       BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::100/120           23          93        BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
```

When DHCPv6 relay binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 relay binding detail
Session Id:  6
     Hardware Address:                    00:10:94:00:00:01
     Client IPv6 Address:                 7001:2:3::d/128
     Lease Expires:                       2017-12-11 07:45:27 IST
     Lease Expires in:                    9999952 seconds
     Preferred Lease Expires:             2017-12-11 07:45:27 IST
     Preferred Lease Expires in:          9999952 seconds
     Client IPv6 Prefix:                  7001::1000:0:0:0/68
     Client IPv6 Excluded Prefix:         7001::1fff:ffff:ffff:ff00/120
     Lease Expires:                       2017-12-11 07:45:27 IST
     Lease Expires in:                    9999952 seconds
     Preferred Lease Expires:             2017-12-11 07:45:27 IST
     Preferred Lease Expires in:          9999952 seconds
     Client DUID:                         LL_TIME0x1-0x599553b0-00:10:94:00:00:01
     State:                               BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Start:                         2017-08-17 13:58:33 IST
     Last Packet Received:                2017-08-17 13:58:48 IST
     Incoming Client Interface:           ge-0/0/0.100
     Client Interface Vlan Id:            100
     Server Ip Address:                   7002::1
     Server Interface:                    none
     Client Id Length:                    14
     Client Id:                           /0x00010001/0x599553b0/0x00109400/0x0001
     Generated Circuit ID:                ge-0/0/0:100
```

show dhcpv6 relay binding detail (Subscriber with Multiple Addresses)

```
user@host> show dhcpv6 relay binding detail
Session Id:  3
```

```
    Client IPv6 Address:                2001:db8:1001::1:2/128
    Lease Expires:                      2015-05-15 02:34:51 PDT
    Lease Expires in:                   24 seconds
    Preferred Lease Expires:            2015-05-15 02:34:51 PDT
    Preferred Lease Expires in:         24 seconds
    Client IPv6 Address:                2001:db8:1001::1:12/128
    Lease Expires:                      2015-05-15 02:41:31 PDT
    Lease Expires in:                   424 seconds
    Preferred Lease Expires:            2015-05-15 02:41:31 PDT
    Preferred Lease Expires in:         424 seconds
    Client IPv6 Address:                2001:db8:1001::1:a/128
    Lease Expires:                      2015-05-15 02:38:11 PDT
    Lease Expires in:                   224 seconds
    Preferred Lease Expires:            2015-05-15 02:38:11 PDT
    Preferred Lease Expires in:         224 seconds
    Client IPv6 Prefix:                 2001:db8:3001::/120
    Lease Expires:                      2015-05-15 02:34:51 PDT
    Lease Expires in:                   24 seconds
    Preferred Lease Expires:            2015-05-15 02:34:51 PDT
    Preferred Lease Expires in:         24 seconds
    Client IPv6 Prefix:                 2001:db8:3001::200/120
    Lease Expires:                      2015-05-15 02:38:11 PDT
    Lease Expires in:                   224 seconds
    Preferred Lease Expires:            2015-05-15 02:38:11 PDT
    Preferred Lease Expires in:         224 seconds
    Client IPv6 Prefix:                 2001:db8:3001::100/120
    Lease Expires:                      2015-05-15 02:36:31 PDT
    Lease Expires in:                   124 seconds
    Preferred Lease Expires:            2015-05-15 02:36:31 PDT
    Preferred Lease Expires in:         124 seconds
    Client DUID:                        LL_TIME0x1-0x55554c6e-00:10:94:00:00:02
    State:                              BOUND(DHCPV6_RELAY_STATE_BOUND)
    Lease Start:                        2015-05-15 02:34:21 PDT
    Last Packet Received:               2015-05-15 02:34:22 PDT
    Incoming Client Interface:          ge-9/0/9.0
    Client Interface Vlan Id:           111
    Demux Interface:                    demux0.3221225475
    Server Ip Address:                  2001:db8:5001::1
    Server Interface:                   none
    Client Profile Name:                DHCP-IPDEMUX-PROF
    Client Id Length:                   14
    Client Id:                          /0x00010001/0x55554c6e/0x00109400/0x0002
    Generated Circuit ID:               ge-9/0/9:111
```

```
    Generated Remote ID Enterprise Number: 1411
    Generated Remote ID:                   ge-9/0/9:111
```

show dhcpv6 relay binding (Interfaces VLAN)

```
user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:DB8::/32        11        87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:DB8::/32        30       79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:DB8::/32        22       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:C9::/32     24       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding summary

```
user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

Release Information

Command introduced in Junos OS Release 11.4.

interfaces-vlan and *interfaces-wildcard* options introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings | 896

clear dhcpv6 relay binding | 2576

show dhcpv6 relay lockout-entries

IN THIS SECTION

- Syntax | 2850
- Description | 2851
- Options | 2851
- Required Privilege Level | 2851
- Output Fields | 2851
- Sample Output | 2853
- Release Information | 2853

Syntax

```
show dhcpv6 relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 relay agent lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 121 on page 2851 lists the output fields for the `show dhcpv6 relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 121: show dhcpv6 relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`

Table 121: show dhcpv6 relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Type of lockout period for the entry: • Grace—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • Lockout—Client is currently locked out; attempts to establish a session are rejected.	all and index
Expires (s)	Number of seconds until the current lockout period expires.	all only
Elapsed (s)	Number of seconds since the current lockout or grace timer started.	all only
Count	Number of consecutive times the client has been locked out.	all only
Expires	Date and time when the current lockout period ends.	index only
Expires in	Number of seconds until the current period expires.	index only
Lockout count	Number of consecutive times client has been locked out.	index only
Next lockout time	Duration of the next lockout period for this client.	index only
Min lockout time	Minimum duration for a lockout period; the initial lockout time.	index only

Table 121: show dhcpv6 relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcpv6 relay lockout-entries (All Entries)

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcpv6 relay lockout-entries (Specific Entry)

```
user@host> show dhcpv6 relay lockout-entries index 2
     Index:                              2
     Key:                                default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                              Lockout
     Expires:                            2018-03-29 19:06:17 IST
     Expires in:                         87
     Lockout count:                      1
     Next lockout time:                  200
     Min lockout time:                   100
     Lockout reason:                     181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

clear dhcpv6 relay lockout-entries | 2580
show dhcp relay lockout-entries | 2796
show dhcp server lockout-entries | 2821
show dhcpv6 server lockout-entries | 2874

show dhcpv6 relay statistics

IN THIS SECTION

- Syntax | 2854
- Description | 2854
- Options | 2855
- Required Privilege Level | 2855
- Output Fields | 2855
- Sample Output | 2859
- Release Information | 2860

Syntax

```
show dhcpv6 relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

bulk-leasequery-connections	(Optional) Display DHCPv6 relay bulk leasequery statistics.
leasequery	(Optional) Display information about DHCPv6 relay individual leasequery statistics.
logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 122 on page 2856 lists the output fields for the `show dhcpv6 relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 122: show dhcpv6 relay statistics Output Fields

Field Name	Field Description
DHCPv6 Packets dropped	Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • Total—Total number of packets discarded by the DHCPV6 relay agent application. • Bad options—Number of packets discarded because invalid options were specified. • Bad send—Number of packets that the extended DHCP relay application could not send. • Bad src address—Number of packets discarded because the family type was not AF_INET6. • Client MAC validation—Number of packets discarded because validation of the client MAC address failed. • No client id—Number of packets discarded because they could not be matched to a client. • Lease Time Violation—Number of packets discarded because of a lease time violation • No safd—Number of packets discarded because they arrived on an unconfigured interface. • Short packet—Number of packets discarded because they were too short. • Relay hop count—Number of packets discarded because the hop count in the packet exceeded 32.

Table 122: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
Messages received	Number of DHCPv6 messages received. • DHCPv6_DECLINE—Number of DHCPv6 PDUs of type DECLINE received • DHCPV6_SOLICIT—Number of DHCPv6 PDUs of type SOLICIT received • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received • DHCPV6_RELEASE—Number of DHCPv6 PDUs of type RELEASE received • DHCPV6_REQUEST—Number of DHCPv6 PDUs of type REQUEST received • DHCPV6_CONFIRM—Number of DHCPv6 PDUs of type CONFIRM received DHCPV6_RENEW—Number of DHCPv6 PDUs of type RENEW received DHCPV6_REBIND—Number of DHCPv6 PDUs of type REBIND received DHCPV6_RELAY_REPL—Number of DHCPv6 PDUs of type RELAY-REPL received • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs received. • DHCPV6_LEASEQUERY_REPLY—Number of DHCPv6 replies received from the DHCPv6 sever • DHCPV6_LEASEQUERY_DATA—xxxx • DHCPV6_LEASEQUERY_DONE—The leasequery is complete
Messages sent	Number of DHCPv6 messages sent. • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted • DHCP_REPLY—Number of DHCPv6 REPLY PDUs transmitted • DHCP_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted • DHCP_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs transmitted • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs transmitted. • DHCP6_LEASEQUERY—Number of DHCP leasequery messages transmitted

Table 122: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
Packets forwarded	Number of packets forwarded by the extended DHCPv6 relay agent application. • FWD REQUEST—Number of DHCPv6 REQUEST packets forwarded • FWD REPLY—Number of DHCPv6 REPLY packets forwarded
External Server Response	State of the external DHCP server responsiveness.
Total Requested Servers	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
Total Attempted Servers	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.
Total Connected	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.
Total Terminated by Server	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
Total Max Attempted	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
Total Closed due to Errors	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
In-Flight Connected	Number of current bulk leasequery connections on the DHCP relay agent.
Bulk Leaasequery Reply Packet Retries	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

Sample Output

show dhcpv6 relay statistics

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       2
    Lease Time Violation        1
    Client MAC validation       1

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_FORW           0
    DHCPV6_LEASEQUERY_REPLY     0
    DHCPV6_LEASEQUERY_DATA      0
    DHCPV6_LEASEQUERY_DONE      0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_REPL           0
    DHCPV6_LEASEQUERY           0

Packets forwarded:
    Total                       4
    FWD REQUEST                 2
    FWD REPLY                   2

External Server Response:
    State                     Responding
```

show dhcpv6 relay statistics bulk-leasequery-connections

```
user@host> show dhcpv6 relay statistics bulk-leasequery-connections

Total Requested Servers:    0
Total Attempted Servers:    0
Total Connected:            0
Total Terminated by Server: 0
Total Max Attempted:        0
Total Closed due to Errors: 0
In-Flight Connected:        0
Bulk Leasequery Reply Packet Retries:   0
```

Release Information

Command introduced in Junos OS Release 11.4.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcpv6 relay statistics | 2583

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | 923

show dhcpv6 server active-leasequery

IN THIS SECTION

- Syntax | 2861
- Description | 2861
- Required Privilege Level | 2861
- Output Fields | 2861
- Sample Output | 2861
- Release Information | 2862

Syntax

```
show dhcpv6 server active-leasequery summary
```

Description

Display the active leasequery status summary of the DHCPv6 local server.

Required Privilege Level

view

Output Fields

Table 123 on page 2861 lists the output fields for the `show dhcpv6 server active-leasequery summary` command.

Table 123: show dhcpv6 server active-leasequery summary Output Fields

Field Name	Field Description
`Connected Peers`	Number of peer DHCP servers connected.
`Connecting Peers`	Number of DHCP server peers trying to connect or reconnect after failure.
`peer`	IP address of the peer DHCPserver.

Sample Output

show dhcp server active-leasequery summary

```
user@host> show dhcp server active-leasequery summary
```

```
            Active Lease-query peer List:
```

```
peer                Connected Peers         Connecting Peers
192.0.2.0           1                       0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

M:N Subscriber Service Redundancy on DHCP Server | 1219

show dhcpv6 server active-leasequery statistics

IN THIS SECTION

- Syntax | 2862
- Description | 2862
- Options | 2863
- Required Privilege Level | 2863
- Output Fields | 2863
- Sample Output | 2864
- Release Information | 2864

Syntax

```
show dhcpv6 server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Display the active leasequery statistics of the DHCPv6 local server.

Options

peer *peer-address*	IP address of the peer DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

view

Output Fields

Table 124 on page 2863 lists the output fields for the `show dhcpv6 server active-leasequery statistics` command.

Table 124: show dhcpv6 server active-leasequery statistics Output Fields

Field Name	Field Description
`ALQ Transmit Buffer count`	Numbers of packets the active leasequery transmitted.
`Bindings Failed to install`	Number of binding installation attempt failed.
`Bindings Installed Successfully`	Number of successful binding installation.
`Bindings Received`	Number of binding request received.
`Bindings Sent`	Number of binding request sent.
`Last Synchronization Time`	Last synchronization time.
`Local Interface count`	Number of local interfaces.
`Max Leasequery Transmit Rate`	Maximum leasequery transmit rate.

Table 124: show dhcpv6 server active-leasequery statistics Output Fields *(Continued)*

Field Name	Field Description
peer	IP address of the peer DHCP server.
Remote Interface count	Number of remote interfaces.
State	Status of the topology discovery configuration.
Topology-Discover Configured	Status of the topology discovery.

Sample Output

show dhcpv6 server active-leasequery peer ip-address

```
user@host> show dhcpv6 server active-leasequery peer 192.0.2.0
```

```
    peer : 192.0.2.0
    Topology-Discover Configured          : Yes
    State                                 : Done
    Bindings Sent                         : 0
    Bindings Received                     : 4
    Bindings Installed Successfully       : 0
    Bindings Failed to install            : 0
    Last Synchronization Time             : 2021-11-22 17:53:52 IST
    ALQ Transmit Buffer count             : 65535
    Max Leasequery Transmit Rate          : 60
    Local Interface count                 : 1
    Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear dhcpv6 server active-leasequery statistics | 2585

show dhcpv6 server binding

IN THIS SECTION

- Syntax | 2865
- Description | 2865
- Options | 2866
- Required Privilege Level | 2866
- Output Fields | 2866
- Sample Output | 2869
- Release Information | 2873

Syntax

```
show dhcpv6 server binding
<address>
<brief | detail | summary>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.

Options

address	(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show: • *CID*—The specified Client ID (CID). • *ipv6-prefix*—The specified IPv6 prefix. • *session-id*—The specified session ID.
brief | detail | summary	(Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcpv6 server binding`.
interface *interface-name*	(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.
interfaces-vlan	(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.
interfaces-wildcard	(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 125 on page 2867 lists the output fields for the `show dhcpv6 server binding` command. Output fields are listed in the approximate order in which they appear.

Table 125: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCPv6 local server: • BOUND—Client has active IP address lease. • INIT—Initial state. • RECONFIGURE—Server has sent reconfigure message to client. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers.	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief
Client IPv6 Address	Client's IPv6 address.	detail

Table 125: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Client IPv6 Prefix	Client's IPv6 prefix.	detail
Client IPv6 Excluded Prefix	IPv6 Prefix of the DHCP client excluded.	detail
Client DUID	Client's DHCP Unique Identifier (DUID).	brief detail
Lease expires	Date and time at which the client's IP address lease expires.	detail
Lease expires in	Number of seconds in which lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail

Table 125: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail
Server Id	DHCP unique identifier (DUID) for the DHCPv6 server.	detail
Client Interface Svlan Id	S-VLAN ID of the client's incoming interface.	detail
Client Interface Vlan Id	VLAN ID of the client's incoming interface.	detail
Dual Stack Group	DHCPv6 server profile name.	detail
Dual Stack Peer Address	DHCPv6 Peer IP address.	detail

Sample Output

show dhcpv6 server binding

```
user@host> show dhcpv6 server binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1111:2222::/64 6           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:db8:1111:2222::/64 7           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:db8:1111:2222::/64 8           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
```

```
2001:db8:1111:2222::/64 9           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:db8:1111:2222::/64 10          86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2001:db8:2002::1/74 11              86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06
```

show dhcpv6 server binding detail

```
user@host> show dhcpv6 server binding detail
Session Id:  2
     Client IPv6 Prefix:                   2001:db8:ffff:0:4::/64
     Client IPv6 Address:                  2001:db8:0:8003::1/128
     Client DUID:                          LL0x1-00:00:64:01:01:02
     State:                                BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Expires:                        2016-11-07 08:30:39 PST
     Lease Expires in:                     43706 seconds
     Preferred Lease Expires:              2016-11-07 08:30:39 PST
     Preferred Lease Expires in:           43706 seconds
     Lease Start:                          2016-11-04 11:00:37 PDT
     Last Packet Received:                 2016-11-06 09:00:39 PST
     Incoming Client Interface:            ae0.3221225472
     Client Interface Svlan Id:            2000
     Client Interface Vlan Id:             1
     Server Ip Address:                    2001:db8::2
     Server Interface:                     none
     Client Profile Name:                  my-dual-stack
     Client Id Length:                     10
     Client Id:                            /0x00030001/0x00006401/0x0102
     Dual Stack Group:                     my-dual-stack
     Dual Stack Peer Address:              192.0.2.10
```

command-name

When DHCPv6 binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 server binding detail
Session Id:  5
     Client IPv6 Address:                  2001:db8:2:3::d/128
     Lease Expires:                        2017-12-11 07:45:15 IST
```

```
     Lease Expires in:                    9999995 seconds
     Preferred Lease Expires:             2017-12-11 07:45:15 IST
     Preferred Lease Expires in:          9999995 seconds
     Client IPv6 Prefix:                  2001:db8::1000:0:0/68
       Client IPv6 Excluded Prefix:         2001:db8::1fff:ffff:ff00/120
     Lease Expires:                       2017-12-11 07:45:15 IST
     Lease Expires in:                    9999995 seconds
     Preferred Lease Expires:             2017-12-11 07:45:15 IST
     Preferred Lease Expires in:          9999995 seconds
     Client DUID:                         LL_TIME0x1-0x599553b0-00:10:94:00:00:01
     State:                               BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Start:                         2017-08-17 13:58:32 IST
     Last Packet Received:                2017-08-17 13:58:36 IST
     Incoming Client Interface:           ge-0/0/0.0
     Client Interface Vlan Id:            100
     Client Pool Name:                    ia_na_pool
     Client Prefix Pool Name:             prefix_delegate_pool
     Client Id Length:                    14
     Client Id:                           /0x00010001/0x599553b0/0x00109400/0x0001
     Relay Id Length:                     31
     Relay Id:                            /0x00020000/0x05830130/0x303a3035/0x3a38363a
     Relay Id:                            /0x34343a65/0x323a6330/0x00000000/0x000000
```

show dhcpv6 server binding interface

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1111:2222::/64 1           86055    BOUND    ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01
```

show dhcpv6 server binding interface detail

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id:  7
     Client IPv6 Prefix:                  2001:db8:1111:2222::/64
     Client DUID:                         LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                               BOUND(bound)
     Lease Expires:                       2009-07-21 10:41:15 PDT
     Lease Expires in:                    86136 seconds
     Preferred Lease Expires:             2012-07-24 00:18:14 UTC
```

```
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (IPv6 Prefix)

```
user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005 detail
Session Id:  7
     Client IPv6 Prefix:                    2001:db8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86136 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (Session ID)

```
user@host> show dhcpv6 server binding 8
Prefix           Session Id  Expires  State    Interface    Client DUID
2001:db8::/32    8           86235    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
```

show dhcpv6 server binding (Interfaces VLAN)

```
user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:db8::/32     11          87583    BOUND    ge-1/0/0.1073741827
```

```
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:db8::/32     30          79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:db8::/32     22          79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     24       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

Release Information

Command introduced in Junos OS Release 9.6.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings | 896

clear dhcpv6 server binding | 2587

show dhcpv6 server lockout-entries

IN THIS SECTION

- Syntax | 2874
- Description | 2874
- Options | 2874
- Required Privilege Level | 2875
- Output Fields | 2875
- Sample Output | 2876
- Release Information | 2877

Syntax

```
show dhcpv6 server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 local server lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 126 on page 2875 lists the output fields for the `show dhcpv6 server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 126: show dhcpv6 server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only

Table 126: show dhcpv6 server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Expires	Date and time when the current lockout period ends.	index only
Expires in	Number of seconds until the current period expires.	index only
Lockout count	Number of consecutive times client has been locked out.	index only
Next lockout time	Duration of the next lockout period for this client.	index only
Min lockout time	Minimum duration for a lockout period; the initial lockout time.	index only
Lockout reason	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	index only

Sample Output

show dhcpv6 server lockout-entries (All Entries)

```
user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200           2
2     00:00:5E:00:53:11     Grace      120           780            2
3     00:00:5E:00:53:22     Lockout    180           2300           1
```

show dhcpv6 server lockout-entries (Specific Entry)

```
user@host> show dhcpv6 server lockout-entries index 2
     Index:                              2
     Key:                                default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                              Lockout
```

```
    Expires:                          2018-03-29 19:06:17 IST
    Expires in:                       87
    Lockout count:                    1
    Next lockout time:                200
    Min lockout time:                 100
    Lockout reason:                   181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- clear dhcpv6 server lockout-entries | 2591
- show dhcp relay lockout-entries | 2796
- show dhcp server lockout-entries | 2821
- show dhcpv6 relay lockout-entries | 2850

show dhcpv6 server statistics

IN THIS SECTION

- Syntax | 2878
- Description | 2878
- Options | 2878
- Required Privilege Level | 2878
- Output Fields | 2878
- Sample Output | 2881
- Release Information | 2882

Syntax

```
show dhcpv6 server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCPv6 local server bulk leasequery statistics.
logical-system *logical-system-name*	(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 127 on page 2879 lists the output fields for the `show dhcpv6 server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 127: show dhcpv6 server statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • `Total`—Total number of packets discarded by the extended DHCPv6 local server • `Strict Reconfigure`—Number of solicit messages discarded because the client does not support reconfiguration • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified • `Bad opcode`—Number of packets discarded because an invalid operation code was specified • `Bad options`—Number of packets discarded because invalid options were specified • `Client MAC validation`—Number of packets discarded because validation of the client MAC address failed. • `Invalid server address`—Number of packets discarded because an invalid server address was specified • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment • `No interface match`—Number of packets discarded because they did not belong to a configured interface • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance • `No valid local address`—Number of packets discarded because there was no valid local address • `Packet too short`—Number of packets discarded because they were too short • `Read error`—Number of packets discarded because of a system read error • `Send error`—Number of packets that the extended DHCPv6 local server could not send

Table 127: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
Advertise Delay	Number of DHCP advertise messages delayed. • DELAYED—Number of DHCPv6 advertise packets that have been sent after being delayed. • INPROGRESS—Number of DHCPv6 advertise packets that are in the delay queue. • TOTAL—Total number of delayed DHCPv6 advertise messages; sum of DELAYED and INPROGRESS.
Messages received	Number of DHCPv6 messages received. • DHCPV6_CONFIRM—Number of DHCPv6 CONFIRM PDUs received. • DHCPV6_DECLINE—Number of DHCPv6 DECLINE PDUs received. • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • DHCPV6_REBIND—Number of DHCPv6 REBIND PDUs received. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs received. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs received. • DHCPV6_RELEASE—Number of DHCPv6 RELEASE PDUs received. • DHCPV6_RENEW—Number of DHCPv6 RENEW PDUs received. • DHCPV6_REQUEST—Number of DHCPv6 REQUEST PDUs received. • DHCPV6_SOLICIT—Number of DHCPv6 SOLICIT PDUs received. • DHCPV6_LEASEQUERY—Number of DHCPv6 leasequery messages received.

Table 127: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
Messages sent	Number of DHCPv6 messages sent. • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_REPLY—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_LOGICAL_NAK—Number of logical NAK messages sent, signifying T1 and T2 timers with values of zero; subset of DHCPV6_REPLY counter. (Displays only at verbose level. • DHC6_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs transmitted. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs transmitted. • DHCPV6_LEASEQUERY_REPLY—Number of DHCPv6 leasequery replies transmitted to the DHCPv6 relay agent. • DHCPV6_LEASEQUERY_DATA—Number of DHCPv6 LEASEQUERY-DATA packets transmitted. • DHCPV6_LEASEQUERY_DONE—Number of DHCPv6 LEASEQUERY-DONE packets sent.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
Dhcpv6 Packets dropped:
    Total                      2
    Lease Time Violation       1
    Client MAC validation      1

Advertise Delay:
    DELAYED                    3
    INPROGRESS                 9
    TOTAL                     12

Messages received:
    DHCPV6_DECLINE             0
```

```
    DHCPV6_SOLICIT             9
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             5
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
    DHCPV6_RELAY_FORW          0

    DHCPV^_LEASEQUERY          0

Messages sent:
    DHCPV6_ADVERTISE           9
    DHCPV6_REPLY               5
    DHCPV6_RECONFIGURE         0
    DHCPV6_RELAY_REPL          0
    DHCPV6_LEASEQUERY_REPLY    0
    DHCPV6_LEASEQUERY_DATA     0
    DHCPV6_LEASEQUERY_DONE     0
```

show dhcpv6 server statistics bulk-leasequery-connections

```
user@host> show dhcpv6 server statistics bulk-leasequery-connections

Total Accepted Connections:                    0
Total Not-Accepted Connections:                0
Connections Closed due to Errors:              0
Connections Closed due to max-empty-replies:   0
In-flight Connections:                         0
```

Release Information

Command introduced in Junos OS Release 9.6.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcpv6 server statistics | 2593

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking | 923

show diameter

IN THIS SECTION

- Syntax | 2883
- Description | 2883
- Options | 2883
- Required Privilege Level | 2884
- Output Fields | 2884
- Sample Output | 2889
- Release Information | 2891

Syntax

```
show diameter
<brief | detail | summary>
```

Description

Display information about the Diameter node.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes Diameter node status. The `brief` output adds summary information about functions, instances, network elements, and peers. The `detail` output adds summary information about routes.

Required Privilege Level

view

Output Fields

Table 128 on page 2884 lists the output fields for the `show diameter` command. Output fields are listed in the approximate order in which they appear.

Table 128: show diameter Output Fields

Field Name	Field Description	Level of Output
`Diameter process id`	ID number of the Diameter process.	All levels
`Functions`	Number of functions associated with Diameter.	All levels
`Connected functions`	Number of functions with active Diameter connections.	All levels
`Instances`	Number of configured Diameter instances.	All levels
`Network elements (NEs)`	Number of configured Diameter network elements.	All levels
`Connected NEs`	Number of Diameter network elements with active connections.	All levels
`Peers`	Number of Diameter peer nodes.	All levels
`Activated peers`	Number of Diameter peers with active connections.	All levels
`Open peers`	Number of peers in the open state, without active network element connections but available for a connection.	All levels
`Transports`	Number of transports configured.	All levels

Table 128: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Requests queued for network transmit`	Number of requests waiting to be sent to the Diameter peers.	All levels
`Answers queued for network transmit`	Number of replies waiting to be sent to the Diameter peers.	All levels
`Expected answers from network`	Number of replies expected to be received from the Diameter peers.	All levels
`Requests queued for function transmit`	Number of requests waiting to be sent to the functions associated with Diameter.	All levels
`Answers queued for function transmit`	Number of replies waiting to be sent to the functions associated with Diameter.	All levels
`Expected answers from functions`	Number of replies expected to be received from the functions associated with Diameter.	All levels
`Memory used by network transmit queues`	Amount of memory consumed by network transmit queues.	All levels
`Memory used by function transmit queues`	Amount of memory consumed by function transmit queues.	All levels
`Origin-state-id`	Value of the Origin-State-ID AVP.	All levels
`Function`	Name of the function for which information is displayed.	`briefdetail`

Table 128: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the Diameter connection with the function: Connected or Disconnec (disconnected).	briefdetail
Upstream Transaction Utilization	Percent of upstream traffic used for this function.	briefdetail
Downstream Transaction Utilization	Percent of downstream traffic used for this function.	briefdetail
Net Queue Buffer Utilization	Percent of network transmission buffer used for this function.	briefdetail
Func Queue Buffer Utilization	Percent of function transmission buffer used for this function.	briefdetail
Routed Dests	Number of destinations that have this function associated with their routes.	briefdetail
Name	Name of the Diameter instance.	briefdetail
Origin-Realm	Value of Origin-Realm attribute-value pair (AVP).	briefdetail
Origin Host	Value of Origin-Host AVP.	briefdetail
NE-Total	Number of configured network elements.	briefdetail
NE-Connected	Number of network elements with active Diameter connections.	briefdetail
Name	Name of the Diameter network element.	briefdetail

Table 128: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Instance`	Name of the Diameter instance in which the network element is configured.	`briefdetail`
`State`	State of the network element: • Connecting—None of the network element peers are in the open state and available for connection. • Selecting—One network element peer is connected and the network element is waiting for another peer to reach the open state so that it can be connected. • Partially-Connected—One network element peer is in the open state and connected. • Post-selection-delay—Three or more peers are in the open state and the network element is waiting to deactivate the peers in excess of two. • Fully-connected—Two network element peers are in the open state and connected.	`briefdetail`
`Primary Peer`	Primary peer for the network element, based on the configured peer priority.	`briefdetail`
`Secondary Peer`	Secondary peer for the network element, based on the configured peer priority.	`briefdetail`
`Peer`	Name of the peer.	`briefdetail`
`Instance`	Name of the Diameter instance in which the peer is configured.	`briefdetail`

Table 128: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the peer: • Bad-Config—Misconfiguration. • Bad-Remote—Remote side does not conform to one of the decisions or is sending malformed messages. • Closed—Normal disconnect due to a request from the remote site or due to excessive watchdog timeouts. • Destructing—Peer to be deleted on the next timer tick. Until then, it performs no actions. • Disabled—Peer is administratively disabled. • Internal-error—Internal error has been detected and the peer is in the process of restarting. • No-Activation—Peer is not used by any Diameter network element. • Rejected—Connection was rejected by remote side of the connection. • Suspended—All other reasons to be suspended.	briefdetail
NE-Count	Number of network elements associated with the peer.	briefdetail
Activated Count	Activation status of the peer: • 1—Peer is activated. • 0—Peer is not activated.	briefdetail
Primary Count	Status of the peer: primary (1) or secondary (0).	briefdetail
Secondary Count	Status of the peer: secondary (0) or primary (1).	briefdetail
Route	Name of the Diameter route.	detail

Table 128: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
NE	Name of the Diameter network element in which the route is configured	detail
Instance	Name of the Diameter instance in which the route is configured.	detail
Valid	Determination of whether the route is valid: yes or no.	detail
Up	State of the route: yes for an active route, no for an inactive route.	detail

Sample Output

show diameter brief

```
user@host> show diameter brief

Diameter node:
  Diameter process id                    :      1446
  Functions                              :   4
  Connected functions                    :   2
  Instances                              :   1
  Network elements(NEs)                  :   1
  Connected NEs                          :   0
  Peers                                  :   2
  Activated peers                        :   1
  Open peers                             :   0
  Transports                             :   1
  Requests queued for network transmit   :         0
  Answers queued for network transmit    :         0
  Expected answers from network          :         0
  Requests queued for function transmit  :         0
  Answers queued for function transmit   :         0
  Expected answers from functions        :         0
  Memory used by network transmit queues :         0
  Memory used by function transmit queues :        0
```

```
  Origin-state-id                        :         0

Diameter function list:
                     Upstream    Downstream  Net Queue   Func Queue
                     Transaction Transaction Buffer      Buffer      Routed
                     Utilization Utilization Utilization Utilization Dests
  Function  State    %           %           %           %
  charging- Disconnec   0           0           0           0           0
  gx-plus   Connected   0           0           0           0           1
  jsrc      Connected   0           0           0           0           0
  packet-tr Disconnec   0           0           0           0           0

Diameter instances:
  Name         Origin-Realm         Origin-Host          NE-Total  NE-Connected
  master       orrr                 ohhh                        1         0

Diameter network-elements:
                                          Primary      Secondary
  Name         Instance     State           Peer         Peer
  n0           master       Connecting      <NONE>       <NONE>

Diameter peer list:
                                                    Activated Primary Secondary
  Peer         Instance     State          NE-Count Count     Count   Count
  p0           master       Suspended        1        1         0       0
  p100         master       No-Activation    0        0         0       0
```

show diameter detail

```
user@host> show diameter detail
...
Diameter routes:
  Route        NE           Instance     Valid Up
  dne-route1   dne1         master       yes   no
```

show diameter summary

```
user@host> show diameter summary
Diameter node:
  Diameter process id                    :      1446
```

```
Functions                              :   4
Connected functions                    :   2
Instances                              :   1
Network elements(NEs)                  :   1
Connected NEs                          :   0
Peers                                  :   2
Activated peers                        :   1
Open peers                             :   0
Transports                             :   1
Requests queued for network transmit   :         0
Answers queued for network transmit    :         0
Expected answers from network          :         0
Requests queued for function transmit  :         0
Answers queued for function transmit   :         0
Expected answers from functions        :         0
Memory used by network transmit queues :         0
Memory used by function transmit queues :        0
Origin-state-id                        :         0
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- clear diameter function statistics | 2595
- clear diameter peer | 2597
- show diameter function | 2892
- show diameter instance | 2903
- show diameter network-element | 2905
- show diameter peer | 2914
- show diameter route | 2929

show diameter function

IN THIS SECTION

- Syntax | 2892
- Description | 2892
- Options | 2892
- Required Privilege Level | 2893
- Output Fields | 2893
- Sample Output | 2895
- Release Information | 2897

Syntax

```
show diameter function
<brief | detail | summary>
<function-name>
```

Description

Display information about all functions associated with Diameter instances or only the specified function.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

function-name (Optional) Display information for only the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions.

Required Privilege Level

view

Output Fields

Table 129 on page 2893 lists the output fields for the `show diameter function` command. Output fields are listed in the approximate order in which they appear.

Table 129: show diameter function Output Fields

Field Name	Field Description	Level of Output
`Function name`	Name of the function for which information is displayed.	All levels
`State`	State of the Diameter connection with the function.	All levels
`Upstream transaction utilization`	Percent of upstream traffic used for this function.	All levels
`Downstream transaction utilization`	Percent of downstream traffic used for this function.	All levels
`Network transmit buffer utilization`	Percent of network transmission buffer used for this function.	All levels
`Function transmit buffer utilization`	Percent of function transmission buffer used for this function.	All levels
`Routed destinations`	Number of destinations that have this function associated with their routes.	All levels

Table 129: show diameter function Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Requests queued for network tx`	Number of requests waiting to be sent to the Diameter peers for this function.	`detail`
`Pending answers from network`	Number of replies expected from the Diameter peers for this function.	`detail`
`Answers queued for function tx`	Number of replies waiting to be sent to this function.	`detail`
`Total upstream transactions pending`	Total number of messages queued for this function.	`detail`
`Upstream transactions limit`	Total number of messages queued for this function.	`detail`
`Requests queued for function tx`	Number of requests waiting to be sent to this function.	`detail`
`Pending answers from function`	Number of replies expected to be received from this function.	`detail`
`Answers queued for network tx`	Number of replies waiting to be sent to this function.	`detail`
`Total downstream transactions pending`	Total number of messages queued for the Diameter peers.	`detail`
`Downstream transactions limit`	Maximum number of messages that can be queued for the Diameter peers.	`detail`

Table 129: show diameter function Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Buffers used by network tx queue	Number of buffers used by messages queued for the Diameter peers.	detail
Limit on network tx queue buffers	Maximum buffer capacity available for messages queued for the Diameter peers.	detail
Buffers used by function tx queue	Number of buffers used by messages queued for this function.	detail
Limit on function tx queue buffers	Maximum buffer capacity available for messages queued for this function.	detail
Origin-state-id	Value of the Origin-State-ID AVP.	detail

Sample Output

show diameter function

```
user@host> show diameter function

Diameter function list:
                     Upstream    Downstream  Net Queue   Func Queue
                     Transaction Transaction Buffer      Buffer      Routed
                     Utilization Utilization Utilization Utilization Dests
  Function  State    %           %           %           %
  jsrc      Disconnec   0           0           0           0           0
```

show diameter function brief

```
user@host> show diameter function brief
```

```
Diameter function:
  Function name                          : gx-plus
  State                                  : Connected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         1

  Function name                          : jsrc
  State                                  : Disconnected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         0
```

show diameter function detail (JSRC)

```
user@host> show diameter function detail

Diameter function:
  Function name                          : jsrc
  State                                  : Disconnected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         0
  Requests queued for network tx         :         0
  Pending answers from network           :         0
  Answers queued for function tx         :         0
  Total upstream transactions pending    :         0
  Upstream transactions limit            :      1024
  Requests queued for function tx        :         0
  Pending answers from function          :         0
  Answers queued for network tx          :         0
  Total downstream transactions pending  :         0
  Downstream transactions limit          :      1024
  Buffers used by network tx queue       :         0
```

```
  Limit on network tx queue buffers      :  10485760
  Buffers used by function tx queue      :         0
  Limit on function tx queue buffers     :  10485760
```

show diameter function detail (Gx-Plus)

```
user@host> show diameter function gx-plus detail

Diameter function:
  Function name                          : gx-plus
  State                                  : Connected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         1
  Requests queued for network tx         :         0
  Pending answers from network           :         0
  Answers queued for function tx         :         0
  Total upstream transactions pending    :         0
  Upstream transactions limit            :      1024
  Requests queued for function tx        :         0
  Pending answers from function          :         0
  Answers queued for network tx          :         0
  Total downstream transactions pending  :         0
  Downstream transactions limit          :      1024
  Buffers used by network tx queue       :         0
  Limit on network tx queue buffers      :  10485760
  Buffers used by function tx queue      :         0
  Limit on function tx queue buffers     :  10485760
  Origin-state-id                        :         0
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

clear diameter function statistics | 2595

show diameter | 2883

show diameter function statistics | 2898

show diameter function statistics

IN THIS SECTION

- Syntax | 2898
- Description | 2898
- Options | 2898
- Required Privilege Level | 2899
- Output Fields | 2899
- Sample Output | 2901
- Release Information | 2902

Syntax

```
show diameter function statistics
<brief | detail | summary>
<function-name>
```

Description

Display statistics about all functions associated with Diameter instances or only the specified function.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function statistics. The `brief` output displays the summary

information in a different format and adds numbers accumulated since the Diameter node was started. The `detail` output adds information to the `brief` output.

function-name — (Optional) Display information for only the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions. When you specify a function, the `brief` output is displayed by default, even when you explicitly specify `summary`.

Required Privilege Level

view

Output Fields

Table 130 on page 2899 lists the output fields for the `show diameter function statistics` command. Output fields are listed in the approximate order in which they appear.

Table 130: show diameter function statistics Output Fields

Field Name	Field Description	Level of Output
`Function`	Name of the function for which information is displayed.	All levels
`Delivered Requests`	Number of requests delivered by Diameter to the application.	All levels
`Delivered Answers`	Number of answers delivered by Diameter to the application.	All levels
`Delivered Messages`	Total number of messages delivered by Diameter to the application.	All levels
`Forwarded Requests`	Number of requests sent by Diameter to the network.	All levels
`Forwarded Answers`	Number of answers sent by Diameter to the network.	All levels
`Forwarded Messages`	Number of messages sent by Diameter to the network.	All levels

Table 130: show diameter function statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Function name`	Name of the function for which information is displayed.	All levels
`Over-limit network requests`	Number of requests sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Over-limit network answers`	Number of answers sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Over-limit network messages`	Total number of messages sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Failed to deliver requests`	Number of requests sent by Diameter to its application that were not successfully delivered.	`detail`
`Failed to deliver answers`	Number of answers sent by Diameter to its application that were not successfully delivered.	`detail`
`Failed to deliver messages`	Total number of messages sent by Diameter to its application that were not successfully delivered.	`detail`
`Over-limit function requests`	Number of requests sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Over-limit function answers`	Number of answers sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Over-limit function messages`	Total number of messages sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Failed to forward requests`	Number of requests that were not successfully sent by Diameter to the network.	`detail`

Table 130: show diameter function statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Failed to forward answers`	Number of answers that were not successfully sent by Diameter to the network.	`detail`
`Failed to forward messages`	Total number of messages that were not successfully sent by Diameter to the network.	`detail`

Sample Output

show diameter function statistics

```
user@host> show diameter function statistics
Diameter function statistics:
           Delivered Delivered Delivered Forwarded Forwarded Forwarded
  Function Requests  Answers   Messages  Requests  Answers   Messages
  jsrc             0         0         0         0         0         0
```

show diameter function statistics brief

```
user@host> show diameter function statistics brief

Diameter function statistics:
  Function name               : jsrc
  Delivered requests          :         0         0
  Delivered answers           :         0         0
  Delivered messages          :         0         0
  Forwarded requests          :         0         0
  Forwarded answers           :         0         0
  Forwarded messages          :         0         0
```

show diameter function statistics detail

```
user@host> show diameter function statistics detail

Diameter function statistics:
  Function name                : jsrc
  Delivered requests           :         0         0
  Delivered answers            :         0         0
  Delivered messages           :         0         0
  Forwarded requests           :         0         0
  Forwarded answers            :         0         0
  Forwarded messages           :         0         0
  Over-limit network requests  :         0         0
  Over-limit network answers   :         0         0
  Over-limit network messages  :         0         0
  Failed to deliver requests   :         0         0
  Failed to deliver answers    :         0         0
  Failed to deliver messages   :         0         0
  Over-limit function requests :         0         0
  Over-limit function answers  :         0         0
  Over-limit function messages :         0         0
  Failed to forward requests   :         0         0
  Failed to forward answers    :         0         0
  Failed to forward messages   :         0         0
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

clear diameter function statistics | 2595

show diameter | 2883

show diameter function | 2892

show diameter instance

IN THIS SECTION

- Syntax | 2903
- Description | 2903
- Options | 2903
- Required Privilege Level | 2903
- Output Fields | 2904
- Sample Output | 2904
- Release Information | 2905

Syntax

```
show diameter instance
<brief | detail | summary>
<instance-name>
```

Description

Display information about all Diameter instances or only the specified instance.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic instance information. The `brief` output displays the summary information in a different format. The `detail` output is the same as the `brief` output.

instance-name (Optional) Display information for only the specified Diameter instance.

Required Privilege Level

view

Output Fields

Table 131 on page 2904 lists the output fields for the `show diameter instance` command. Output fields are listed in the approximate order in which they appear.

Table 131: show diameter instance Output Fields

Field Name	Field Description	Level of Output
`name`	Name of the Diameter instance.	`summary`
`Origin-realm`	Value of Origin-Realm AVP.	`summary`
`Origin-host`	Value of Origin-Host AVP.	`summary`
`NE-total`	Total number of network elements configured for this instance.	`summary`
`NE-connected`	Number of network elements with active Diameter connections.	`summary`
`Instance name`	Name of the Diameter instance.	`brief detail`
`Origin realm`	Value of Origin-Realm AVP.	`brief detail`
`Origin host`	Value of Origin-Host AVP.	`brief detail`
`NEs`	Total number of network elements configured for this instance.	`brief detail`
`Connected NEs`	Number of network elements with active Diameter connections.	`brief detail`

Sample Output

show diameter instance

```
user@host> show diameter instance
```

```
Diameter instances:
  Name         Origin-Realm         Origin-Host          NE-Total  NE-Connected
  master       rrrr                 hhhh                        1            1
```

show diameter instance detail

```
user@host> show diameter instance detail

Diameter instance:
  Instance name : master
  Origin realm  : rrrr
  Origin host   : hhhh
  NEs           :   1
  Connected NEs :   1
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 2883

show diameter network-element

IN THIS SECTION

- Syntax | 2906
- Description | 2906
- Options | 2906
- Required Privilege Level | 2906
- Output Fields | 2906

Sample Output | 2908
Release Information | 2909

Syntax

```
show diameter network-element
<brief | detail | summary>
<element-name>
```

Description

Display information about all Diameter network elements or only the specified network element.

Options

brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic network element information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.
element-name	(Optional) Display information for only the specified network element.

Required Privilege Level

view

Output Fields

Table 132 on page 2907 lists the output fields for the `show diameter network-element` command. Output fields are listed in the approximate order in which they appear.

Table 132: show diameter network-element Output Fields

Field Name	Field Description	Level of Output
Name	Name of the Diameter network element.	summary
Instance	Name of the Diameter instance in which the network element is configured.	summary
State	State of the network element: • Connecting—None of the network element peers are in the open state and available for connection. • Selecting—One network element peer is connected and the network element is waiting for another peer to reach the open state so that it can be connected. • Partially-Connected—One network element peer is in the open state and connected. • Post-selection-delay—Three or more peers are in the open state and the network element is waiting to deactivate the peers in excess of two. • Fully-connected—Two network element peers are in the open state and connected.	All levels
Primary peer	Primary peer for the network element, based on the configured peer priority.	All levels
Secondary peer	Secondary peer for the network element, based on the configured peer priority.	All levels
NE name	Name of the Diameter network element.	brief detail
Instance name	Name of the Diameter instance in which the network element is configured.	brief detail
Peers	Number of configured peers.	brief detail

Table 132: show diameter network-element Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Activated peers	Number of peers that have been activated.	brief detail
Open peers	Number of peers in the open state, without active network element connections but available for a connection.	brief detail
Routes	Number of routes configured for the network element.	brief detail
Invalid routes	Number of routes that are invalid because they lack one or more of the following: application and partition, Diameter instance, or destination realm.	brief detail
Activation delay	Period in milliseconds between peer activations by the network element.	brief detail
First selection delay	Period in milliseconds that the network element waited after connecting to the first peer to allow other peers to reach the open state.	brief detail
Post selection delay	Period in milliseconds that the network element waited after having two peers in the open state before deactivating all lower-priority peers.	brief detail

Sample Output

show diameter network-element

```
user@host> show diameter network-element

Diameter network-elements:
                                        Primary      Secondary
  Name         Instance     State          Peer         Peer
  ne0          master       Fully-connected p0          p1
```

show diameter network-element detail

```
user@host> show diameter network-element detail

Diameter network-element:
  NE name               : ne0
  Instance name         : master
  State                 : Fully-connected
  Primary peer          : p0
  Secondary peer        : p1
  Peers                 :   5
  Activated peers       :   4
  Open peers            :   2
  Routes                :   1
  Invalid routes        :   0
  Activation delay      :     10000 ms
  First selection delay :         0 ms
  Post selection delay  :     30000 ms
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- show diameter | 2883
- show diameter function | 2892
- show diameter network-element map | 2910
- show diameter peer | 2914
- show diameter route | 2929

show diameter network-element map

IN THIS SECTION

- Syntax | 2910
- Description | 2910
- Options | 2910
- Required Privilege Level | 2910
- Output Fields | 2911
- Sample Output | 2912
- Release Information | 2913

Syntax

```
show diameter network-element map
<brief | detail | summary>
<element-name>
```

Description

Display network-element-to-peer mapping information for all Diameter network elements or only the specified network element.

Options

brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default. The `brief` output and `detail` output display the summary information in a different format.
element-name	(Optional) Display information for only the specified network element.

Required Privilege Level

view

Output Fields

Table 133 on page 2911 lists the output fields for the `show diameter network-element map` command. Output fields are listed in the approximate order in which they appear.

Table 133: show diameter network-element map Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the Diameter network element.	`summary`
`Instance`	Name of the Diameter instance in which the network element is configured.	`summary`
`Peer`	Name of the peer.	All levels
`Priority`	Priority configured for the peer. A lower number indicates a higher priority.	All levels
`State`	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	`summary`
`NE name`	Name of the Diameter network element.	`brief detail`
`Instance name`	Name of the Diameter instance in which the network element is configured.	`brief detail`

Table 133: show diameter network-element map Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Usage	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	brief detail

Sample Output

show diameter network-element map

```
user@host> show diameter network-element map

Diameter network-element peers:
  Name         Instance     Peer         Priority  State
  ne0          master       p288               30 Activated
  ne0          master       p0                 20 Primary
  ne0          master       pA                 15 Activated
  ne0          master       p1                 10 Secondary
  ne0          master       pB                  5 Not-Activated
```

show diameter network-element map detail

```
user@host> show diameter network-element map detail

Diameter network-element peers:
  NE name               : ne0
  Instance name         : master
  Peer                  : p288
```

```
Priority              :        30
Usage                 : Activated

NE name               : ne0
Instance name         : master
Peer                  : p0
Priority              :        20
Usage                 : Primary

NE name               : ne0
Instance name         : master
Peer                  : pA
Priority              :        15
Usage                 : Activated

NE name               : ne0
Instance name         : master
Peer                  : p1
Priority              :        10
Usage                 : Secondary

NE name               : ne0
Instance name         : master
Peer                  : pB
Priority              :         5
Usage                 : Not-Activated
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 2883

show diameter network-element | 2905

show diameter peer

IN THIS SECTION

- Syntax | 2914
- Description | 2914
- Options | 2914
- Required Privilege Level | 2914
- Output Fields | 2915
- Sample Output | 2918
- Release Information | 2920

Syntax

```
show diameter peer
<brief | detail | summary>
<peer-name>
```

Description

Display information about all peers associated with Diameter instances or only the specified peer.

Options

brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic peer information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.
peer-name	(Optional) Display information for only the specified peer.

Required Privilege Level

view

Output Fields

Table 134 on page 2915 lists the output fields for the `show diameter peer` command. Output fields are listed in the approximate order in which they appear.

Table 134: show diameter peer Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	`brief summary`
`Instance`	Name of the Diameter instance in which the peer is configured.	`brief summary`
`State`	State of the peer: • Bad-Config—Misconfiguration. • Bad-Remote—Remote side does not conform to one of the decisions or is sending malformed messages. • Closed—Normal disconnect due to a request from the remote site or due to excessive watchdog timeouts. • Destructing—Peer to be deleted on the next timer tick; until then, it performs no actions. • Disabled—Peer is administratively disabled. • Internal-error—Internal error has been detected and the peer is in the process of restarting. • No-Activation—Peer is not used by any Diameter network element. • Rejected—Connection was rejected by remote side of the connection. • Reopen–Connection has been unexpectedly closed and Diameter is attempting to reopen the connection. • Suspended—All other reasons to be suspended.	All levels
`NE-Count`	Number of network elements associated with the peer.	`brief summary`

Table 134: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Activated Count	Activation status of the peer: • 1—Peer is activated. • 0—Peer is not activated.	All levels
Primary Count	Status of the peer, primary (1) or secondary (0).	All levels
Secondary Count	Secondary (0) versus Primary (1) status of the peer.	All levels
Peer name	Name of the peer.	detail
NEs	Number of network elements associated with the peer.	detail
Vrf	Logical system:routing instance of the configuration.	detail
Remote address	Remote IP address of the peer.	detail
Remote port	Remote port on the peer on which the connection is made.	detail
Remote end origin realm	Name of the realm of the Diameter node that originates messages to the peer.	detail
Remote end origin host	Name of the host of the Diameter node that originates messages to the peer.	detail
Local address	Local IP address on the Diameter origin node.	detail
Local port	Local port on the Diameter origin node.	detail
Local transport	Number of transports configured.	detail

Table 134: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Time since last enable`	Period since peer was enabled in *hh:mm.ss* format.	`detail`
`In state time`	Period that peer has been in present state in *hh:mm.ss* format.	`detail`
`Remaining in state time`	Period that peer will remain in present state in *hh:mm.ss* format.	`detail`
`Missing wd events`	Number of missed watchdog events.	`detail`
`Tx queue length`	Number of messages in the transmit queue.	`detail`
`Answer waiting count`	Number of answers on which the peer is waiting.	`detail`
`Time since last rx`	Number of milliseconds since the last message was received by the peer.	`detail`
`Time until wd timeout`	Time remaining until next watchdog event.	`detail`
`Operation timeout`	Watchdog timeout period.	`detail`
`Suspended timeout base`	Base timeout period in suspended states (suspended, rejected, bad-remonte, bad-config). This timeout doubles after each consecutive suspension, until the maximum value of 600 seconds is reached.	`detail`
`Closed timeout`	Timeout period in normal closed state, such as when an external peer requested a disconnect.	`detail`
`Connection timeout`	Timeout period for establishing a connection.	`detail`

Table 134: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Waiting origin state id`	Whether the peer is waiting for the Origin-State-Id AVP, `yes` or `no`.	`detail`

Sample Output

show diameter peer

```
user@host> show diameter peer

Diameter peer list:
                                                  Activated Primary Secondary
  Peer          Instance     State          NE-Count Count     Count   Count
  p0            master       I-Open           1        1         1       0
  p1            master       I-Open           1        1         0       1
  p288          master       Suspended        1        1         0       0
  pA            master       Suspended        1        1         0       0
  pB            master       No-Activation    1        0         0       0
  pc            master       No-Activation    0        0         0       0
  pd            master       No-Activation    0        0         0       0
```

show diameter peer detail

```
user@host> show diameter peer detail

Diameter peer:
  Peer name               : p0
  State                   : I-Open
  NEs                     :   1
  Activated count         :   1
  Primary count           :   1
  Secondary count         :   0
  Vrf                     : default:master
  Remote address          : 203.0.113.158
```

```
Remote port             : 62917
Remote end origin realm : rrrrA
Remote end origin host  : hhhhA
Local address           : 203.0.113.155
Local port              : 57095
Local transport         : <NO-TRANSPORT>
Time since last enable  : 08:56.200
In state time           : 08:56.200
Remaining in state time : no limit
Missed wd events        :         0
Tx queue length         :         0
Answer waiting count    :         0
Time since last rx      :      2200 ms
Time until wd timeout   :      3800 ms
Operation timeout       :      6000 ms
Suspended timeout base  :     30000 ms
Closed timeout          :     30000 ms
Connection timeout      :      6000 ms
Waiting origin state id : no

Peer name               : p1
State                   : I-Open
NEs                     :   1
Activated count         :   1
Primary count           :   0
Secondary count         :   1
Vrf                     : default:master
Remote address          : 203.0.113.158
Remote port             : 58490
Remote end origin realm : rrrrA
Remote end origin host  : hhhhB
Local address           : 203.0.113.155
Local port              : 49293
Local transport         : <NO-TRANSPORT>
Time since last enable  : 08:56.200
In state time           : 08:36.000
Remaining in state time : no limit
Missed wd events        :         0
Tx queue length         :         0
Answer waiting count    :         0
Time since last rx      :         0 ms
Time until wd timeout   :      6000 ms
Operation timeout       :      6000 ms
```

```
Suspended timeout base  :     30000 ms
Closed timeout          :     30000 ms
Connection timeout      :      6000 ms
Waiting origin state id : no
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- clear diameter peer | 2597
- show diameter | 2883
- show diameter peer map | 2920
- show diameter peer statistics | 2924

show diameter peer map

IN THIS SECTION

- Syntax | 2921
- Description | 2921
- Options | 2921
- Required Privilege Level | 2921
- Output Fields | 2921
- Sample Output | 2922
- Release Information | 2924

Syntax

```
show diameter peer map
<brief | detail | summary>
<peer-name>
```

Description

Display peer-to-network-element mapping information for all peers associated with Diameter instances or with the specified peer.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic peer information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

peer-name (Optional) Display mapping information for only the specified peer.

Required Privilege Level

view

Output Fields

Table 135 on page 2921 lists the output fields for the `show diameter peer map` command. Output fields are listed in the approximate order in which they appear.

Table 135: show diameter peer map Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	All levels
`Instance`	Name of the Diameter instance in which the network element is configured.	All levels

Table 135: show diameter peer map Output Fields *(Continued)*

Field Name	Field Description	Level of Output
NE	Name of the Diameter network element.	All levels
Priority	Priority configured for the peer. A lower number indicates a higher priority.	All levels
State	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	All levels
Instance name	Name of the Diameter instance in which the network element is configured.	brief detail
NE name	Name of the Diameter network element.	brief detail
Usage	Role of the peer for the network element, Primary or Secondary.	brief detail

Sample Output

show diameter peer map

```
user@host> show diameter peer map

Diameter peer usage by network elements:
  Peer         Instance     NE             Priority State
  p0           master       ne0                  20 Primary
  p1           master       ne0                  10 Secondary
```

```
  p288          master       ne0                  30 Activated
  pA            master       ne0                  15 Activated
  pB            master       ne0                   5 Not-Activated
```

show diameter peer map detail

```
user@host> show diameter peer map detail

Diameter network-element peers:
  Peer                  : p0
  Instance name         : master
  NE name               : ne0
  Priority              :        20
  Usage                 : Primary

  Peer                  : p1
  Instance name         : master
  NE name               : ne0
  Priority              :        10
  Usage                 : Secondary

  Peer                  : p288
  Instance name         : master
  NE name               : ne0
  Priority              :        30
  Usage                 : Activated

  Peer                  : pA
  Instance name         : master
  NE name               : ne0
  Priority              :        15
  Usage                 : Activated

  Peer                  : pB
  Instance name         : master
  NE name               : ne0
  Priority              :         5
  Usage                 : Not-Activated
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

clear diameter peer | 2597
show diameter | 2883
show diameter peer | 2914
show diameter peer statistics | 2924

show diameter peer statistics

IN THIS SECTION

- Syntax | 2924
- Description | 2924
- Options | 2925
- Required Privilege Level | 2925
- Output Fields | 2925
- Sample Output | 2926
- Release Information | 2928

Syntax

```
show diameter peer statistics
<brief | detail | summary>
<peer-name>
```

Description

Display statistics about all peers associated with Diameter instances or only the specified peer.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function statistics. The `brief` output displays the summary information in a different format and adds numbers accumulated since the peer was connected. The `detail` output adds information to the `brief` output.

peer-name (Optional) Display information for only the specified peer. When you specify a peer, the `brief` output is displayed by default, even when you explicitly specify `summary`.

Required Privilege Level

view

Output Fields

Table 136 on page 2925 lists the output fields for the `show diameter peer statistics` command. Output fields are listed in the approximate order in which they appear.

Table 136: show diameter peer statistics Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	`summary brief`
`Instance`	Name of the Diameter instance in which the network element is configured.	`summary brief`
`Rx`	Total number of messages received.	`summary brief`
`Rx-Peer`	Number of messages received by the peer.	`summary brief`
`Rx-node`	Number of messages received by the Diameter node.	`summary brief`
`Forw`	Total number of forwarded messages.	`summary brief`
`Tx-Peer`	Number of messages transmitted by the peer.	`summary brief`

Table 136: show diameter peer statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Tx	Total number of transmitted messages.	summary brief
Peer name	Name of the peer.	detail
Instance name	Name of the Diameter instance in which the network element is configured.	detail

Sample Output

show diameter peer statistics

```
user@host> show diameter peer statistics

Diameter peer statistics:
  Peer         Instance     Rx        Rx-Peer   Rx-Node   Forw      Tx-Peer   Tx
  p0           master            113       113         0         0       113       113
  p1           master            110       110         0         0       110       110
  p288         master              0         0         0         0         0         0
  pA           master              0         0         0         0         0         0
  pB           master              0         0         0         0         0         0
  pc           master              0         0         0         0         0         0
  pd           master              0         0         0         0         0         0
```

show diameter peer statistics detail

```
user@host> show diameter peer statistics detail

Diameter peer statistics:
  Peer name               : p0
  Instance name           : master
                            Current       Since last enable
  Rx errors               :         0             0
```

```
Rx messages             :       114           114
Rx handled by peer      :       114           114
Rx dropped msgs         :         0             0
Rx unmatched answers    :         0             0
Rx answers              :         0             0
Rx requests             :         0             0
Rx total                :         0             0
Forw to connection      :         0             0
Forw to peer            :         0             0
Forw to routed dest     :         0             0
Total forwarding        :         0             0
Forwarding failures     :         0             0
Forwarding success      :         0             0
Moved-in messages       :         0             0
Moved-out messages      :         0             0
Rerouted messages       :         0             0
Dropped tx messages     :         0             0
Tx by peer              :       114           114
Tx errors               :         0             0
Tx total                :       114           114
Connection attempts     :         0             1
Connection fails        :         0             0
Connections             :         0             1
Passive teminations     :         0             0
Active terminations     :         0             0
Passive disconnects     :         0             0
Active disconnects      :         0             0
Rx block requests       :         0             0
Rx block timeoutss      :         0             0
Connection management messages
                          Rx current    Rx since      Tx current    Tx since
                                        last enable                 last enable
CER                     :         0             0             1             1
CEA                     :         1             1             0             0
DWR                     :         0             0           113           113
DWA                     :       113           113             0             0
DPR                     :         0             0             0             0
DPA                     :         0             0             0             0

Peer name               : p1
Instance name           : master
                          Current       Since last enable
Rx errors               :         0             0
```

```
Rx messages             :       110           110
Rx handled by peer      :       110           110
Rx dropped msgs         :         0             0
Rx unmatched answers    :         0             0
Rx answers              :         0             0
Rx requests             :         0             0
Rx total                :         0             0
Forw to connection      :         0             0
Forw to peer            :         0             0
Forw to routed dest     :         0             0
Total forwarding        :         0             0
Forwarding failures     :         0             0
Forwarding success      :         0             0
Moved-in messages       :         0             0
Moved-out messages      :         0             0
Rerouted messages       :         0             0
Dropped tx messages     :         0             0
Tx by peer              :       110           110
Tx errors               :         0             0
Tx total                :       110           110
Connection attempts     :         0             1
Connection fails        :         0             0
Connections             :         0             1
Passive teminations     :         0             0
Active terminations     :         0             0
Passive disconnects     :         0             0
Active disconnects      :         0             0
Rx block requests       :         0             0
Rx block timeoutss      :         0             0
Connection management messages
                         Rx current    Rx since      Tx current    Tx since
                                       last enable                 last enable
CER                     :         0             0             1             1
CEA                     :         1             1             0             0
DWR                     :         0             0           109           109
DWA                     :       109           109             0             0
DPR                     :         0             0             0             0
DPA                     :         0             0             0             0
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

clear diameter peer | 2597
show diameter | 2883
show diameter peer | 2914
show diameter peer map | 2920

show diameter route

IN THIS SECTION

- Syntax | 2929
- Description | 2929
- Options | 2930
- Required Privilege Level | 2930
- Output Fields | 2930
- Sample Output | 2931
- Release Information | 2932

Syntax

```
show diameter route
<brief | detail | summary>
<route-name>
```

Description

Display information about all routes associated with Diameter instances or only the specified route.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

route-name (Optional) Display information for only the specified route.

Required Privilege Level

view

Output Fields

Table 137 on page 2930 lists the output fields for the `show diameter route` command. Output fields are listed in the approximate order in which they appear.

Table 137: show diameter route Output Fields

Field Name	Field Description	Level of Output
`Route`	Name of the route.	`summary brief`
`NE`	Name of the network element associated with the route.	`summary brief`
`Instance`	Name of the Diameter instance in which the route is configured.	`summary brief`
`NE name`	Name of the network element associated with the route.	`brief detail`
`Instance name`	Name of the Diameter instance in which the route is configured.	`brief detail`
`Valid`	Determination whether the route is valid, `yes` or `no`.	All levels
`Up`	State of the route, `yes` (up) or `no` (down).	All levels

Table 137: show diameter route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Function	Name of the function associated with the route.	brief detail
Partition	Partition associated with the function.	brief detail
Dest-realm	Destination realm configured for the route.	brief detail
Dest-host	Destination hostname configured for the route.	brief detail
Metric	Metric associated with the destination and function to create the route.	brief detail
Score	Value that represents how a route is configured. The basic score is 0. Points are added according to the following scheme: • Function is specified—Add 3. • Function partition is specified—Add 1. • Destination realm is specified—Add 1. • Destination host is specified—Add 1.	brief detail

Sample Output

show diameter route

```
user@host> show diameter route

Diameter routes:
  Route        NE           Instance     Valid Up
  rA           ne0          master       yes   yes
```

show diameter route detail

```
user@host> show diameter route detail

Diameter route:
  Route name        : rA
  NE name           : ne0
  Instance mame     : master
  Valid             : yes
  Up                : yes
  Function          : jsrc
  Partition         : jsrc-a
  Dest-realm        : outer-realm
  Dest-host         : outer-host
  Metric            :        50
  Score             :         6
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 2883

show diameter network-element | 2905

show dynamic-profile session

IN THIS SECTION

- Syntax | 2933
- Description | 2933
- Options | 2933
- Required Privilege Level | 2934

Output Fields | 2934
Sample Output | 2934
Release Information | 2938

Syntax

```
show dynamic-profile session
<client-id client-id>
<profile-name profile-name>
<service-id service-id>
```

Description

Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.

NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as `NONE` in the command output.
- The variable is also displayed as `NONE` when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the `show configuration dynamic-profiles` command.

Options

client-id *client-id* Display dynamic profile information for subscribers associated with the specified client.

profile-name *profile-name* (Optional) Display dynamic profile information for the specified subscriber or service profile.

service-id *service-id* Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level

view

Output Fields

This command displays the dynamic client or service profile configuration for each subscriber.

Sample Output

show dynamic-profile session client-id (Client ID)

```
user@host>show dynamic-profile session client-id 20
pppoe {
    interfaces {
        pp0 {
            unit 1073741831 {
                ppp-options {
                    chap;
                    pap;
                }
                pppoe-options {
                    underlying-interface ge-2/0/0.0;
                    server;
                }
                family {
                    inet {
                        unnumbered-address lo0.0;
                    }
                }
            }
        }
    }
    class-of-service {
```

```
        traffic-control-profiles {
            tcp1 {
                scheduler-map smap1_UID1024;
                shaping-rate 100m;
            }
        }
        interfaces {
            pp0 {
                unit 1073741831 {
                    output-traffic-control-profile tcp1;
                }
            }
        }
        scheduler-maps {
            smap1_UID1024 {
                forwarding-class best-effort scheduler sch1_UID1023;
            }
        }
        schedulers {
            sch1_UID1023 {
                transmit-rate percent 40;
                buffer-size percent 40;
                priority low;
            }
        }
    }
}
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
```

```
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
```

```
            }
        }
    }
}
```

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile)

```
user@host>show dynamic-profile session client-id 20 profile-name cos-service
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session service-id (Service Session)

```
user@host>show dynamic-profile session service-id 21
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
```

```
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
```

Release Information

Command introduced in Junos OS Release 13.3.

show ipv6 router-advertisement

IN THIS SECTION

- Syntax | 2939
- Description | 2939
- Options | 2939
- Additional Information | 2940
- Required Privilege Level | 2940
- Output Fields | 2940
- Sample Output | 2942
- Release Information | 2945

Syntax

```
show ipv6 router-advertisement
<conflicts>
<interface interface>
<logical-system (all | logical-system-name)>
<prefix prefix/prefix length>
```

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

The router advertisement module does not function in the backup Routing Engine as the Routing Engine does not send an acknowledgment message after receiving the packets. Starting in Junos OS Release 22.2R1, you can view the router advertisement module information using the `show ipv6 router-advertisement` operational command.

Options

none Display all IPv6 router advertisement information for all interfaces.

conflicts	(Optional) Display only the IPv6 router advertisement information that is conflicting.
interface *interface*	(Optional) Display IPv6 router advertisement information for the specified interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix *prefix/prefix length*	(Optional) Display IPv6 router advertisement information for the specified prefix.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

Output Fields

Table 138 on page 2940 describes the output fields for the `show ipv6 router-advertisement` command. Output fields are listed in the approximate order in which they appear.

Table 138: show ipv6 router-advertisement Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Advertisements sent`	Number of router advertisements sent and the elapsed time since they were sent.
`Solicits received`	Number of solicitation messages received.
`Advertisements received`	Number of router advertisements received.

Table 138: show ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
`Advertisements from`	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
`Managed`	Managed address configuration flag: `0` (stateless) or `1` (stateful).
`Other configuration`	Other stateful configuration flag: `0` (stateless) or `1` (stateful).
`Reachable time`	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
`Default lifetime`	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
`Retransmit timer`	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
`Current hop limit`	Configured current hop limit.
`Prefix`	Name and length of the prefix.
`Valid lifetime`	How long the prefix remains valid for onlink determination.
`Preferred lifetime`	How long the prefix generated by stateless autoconfiguration remains preferred.
`On link`	Onlink flag: `0` (not onlink) or `1` (onlink).
`Autonomous`	Autonomous address configuration flag: `0` (not autonomous) or `1` (autonomous).
`Upstream Mode`	Configured interface as upstream interface for RA proxy
`Downstream Mode`	Configured interface as downstream interface for RA proxy.

Table 138: show ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
Downstream	Downstream interface for RA proxy.
Passive Mode	RA receive only mode is enabled.
Proxy Blackout Timer	Proxy blackout timer interval is the time interval for which the interface must not be used as a proxy interface. Proxy functionality is disabled on that interface.
Parameter Preference	Preference to select configured or proxied parameters for downstream interface
error	Displays the details of the error.

Sample Output

show ipv6 router-advertisement (with RA and NAT64 IPv6 address prefix) (cSRX, MX Series, SRX Series, vMX, and vSRX)

```
user@host> show ipv6 router-advertisement
Interface: ge-0/0/0.0
  Advertisements sent: 3, last sent 00:00:05 ago
  Solicits sent: 1, last sent 00:00:37 ago
  Solicits received: 0
  Advertisements received: 0
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Disable
  Upstream mode: Disable
  Downstream mode: Disable
  Proxy blackout timer: Not Running
  NAT prefix: 6000::/96 Lifetime: 8192 sec
```

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

show ipv6 router-advertisement (Without RA proxy) (SRX Series and vSRX 3.0)

(Without RA proxy)

```
user@host> show ipv6 router-advertisement
Interface: ge-0/0/1.0
  Advertisements sent: 7, last sent 00:00:11 ago
  Solicits sent: 1, last sent 00:00:41 ago
  Solicits received: 0
  Advertisements received: 0
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Enable
  Upstream mode: Disable
  Downstream mode: Disable
  Proxy blackout timer: Not Running
```

(With RA proxy)

```
Interface: ge-0/0/2.0
  Advertisements sent: 2, last sent 00:00:49 ago
```

```
  Solicits sent: 1, last sent 00:01:21 ago
  Solicits received: 0
  Advertisements received: 18
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Enable
  Upstream mode: Enable
  Downstream mode: Disable
  Proxy parameter preference: Proxied
  Proxy blackout timer: Not Running
  Downstream: ge-0/0/0.0
  Downstream: ge-0/0/1.0
  Advertisement from fe80::5668:adff:fed8:101b, heard 00:00:00 ago
  Managed: 0
  Other configuration: 1 [0]
  Link MTU: 1500 bytes
  Reachable time: 5555 ms
  Default lifetime: 1799 sec [1800 sec]
  Retransmit timer: 4444 ms
  Current hop limit: 50 [64]
  RDNSS address: abcd:1::1
  Lifetime: 3333 sec
  Prefix: 2002:2:0:2000::/64
  Valid lifetime: 3600 sec
  Preferred lifetime: 2400 sec
  On link: 1
  Autonomous: 1
  Route Information: 2002:2:0:2000::/64
  IPv6 RA Preference: LOW
  Route lifetime: 1111 sec
  DNSSL suffix: juniper.net
  Lifetime: 6666 sec
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 2001:db8:8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

show ipv6 router-advertisement (Backup Routing Engine)

```
user@host> show ipv6 router-advertisement
error: Module not running (routing)
```

Release Information

Command introduced before Junos OS Release 7.4.

Starting in Junos OS Release 22.1, we support RA proxy on SRX Series and vSRX 3.0 . This command output is modified to display the configured upstream interfaces, downstream interfaces, the proxy flag, the proxy blackout timer, and the passive mode information.

RELATED DOCUMENTATION

clear ipv6 router-advertisement | 2601

show network-access aaa accounting

IN THIS SECTION

- Syntax | 2946
- Description | 2946
- Required Privilege Level | 2946
- Output Fields | 2946
- Sample Output | 2947
- Release Information | 2947

Syntax

```
show network-access aaa accounting
```

Description

Display the state of the RADIUS Acct-On response sent from the RADIUS server.

Required Privilege Level

view

Output Fields

Table 139 on page 2946 lists the output fields for the `show network-access aaa accounting` command. Output fields are listed in the approximate order in which they appear.

Table 139: show network-access aaa accounting Output Fields

Field Name	Field Description
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.

Table 139: show network-access aaa accounting Output Fields *(Continued)*

Field Name	Field Description
Logical System	Logical system associated with the access profile.
Routing Instance	Routing instance associated with the access profile.
Acct-On-Response	Status of the RADIUS Acct-On response. • ACK—ACK response for the Acct-On message is received from the RADIUS server. • ERROR—An error condition has occurred. • NONE— No Acct-On message is sent. • PENDING—Acct-On message is sent to RADIUS server, but no response has been received yet.

Sample Output

show network-access aaa accounting

```
user@host> show network-access aaa accounting
Profile        Logical System   Routing Instance   Acct-On-Response
ppp-profile    default          default            ACK
l2tp-profile   default          l2tp_RI            PENDING
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

RADIUS Acct-On and Acct-Off Messages | 572

show network-access aaa radius-servers

IN THIS SECTION

- Syntax | 2948
- Description | 2948
- Options | 2948
- Required Privilege Level | 2948
- Output Fields | 2948
- Sample Output | 2955
- Release Information | 2958

Syntax

```
show network-access aaa radius-servers
<detail>
```

Description

Display RADIUS server status and information.

Options

detail (Optional) Display detailed level of information.

Required Privilege Level

view

Output Fields

Table 140 on page 2949 lists the output fields for the `show network-access aaa radius-servers` command. Output fields are listed in the approximate order in which they appear.

Table 140: show network-access aaa radius-servers Output Fields

Field Name	Field Description	Level of Output
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.	All levels
`Server address`	IPv4 or IPv6 address of the RADIUS server.	All levels
`Authentication port`	RADIUS server authentication port number.	All levels
`Preauthentication port`	RADIUS server preauthentication port number.	All levels
`Accounting port`	RADIUS server accounting port number.	All levels
`Accounting retry`	Number of times the router retransmits RADIUS accounting messages when no response is received from the server.	Detail
`Accounting timeout`	Period the local router waits to receive a response from a RADIUS accounting server before retransmitting the message.	Detail

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Status	RADIUS server status, UP (Alive), UNREACHABLE, or DOWN (DEAD). If status is DOWN, the Status field includes the number of seconds configured by the revert-interval statement. The router does not send requests to servers in the DOWN state, but does send requests to servers with a status of either UP or UNREACHABLE. This field also displays the status of AAA accounting suspension or resumption, and the status of baselining of accounting statistics if you suspended or resumed accounting operations or initiated the generation of a baseline. This information is applicable only for RADIUS servers that are in the UP state. **NOTE**: After requests to a server or set of servers time out after 10 seconds, the status of the servers changes. The following guidelines apply to server status: • For the purpose of marking a server as Down (DEAD), the request includes the original request and any retries that are configured. The 10-second timeout period starts after the initial request and all retries have expired without receiving a response from the server. The amount of the timeout period that elapses before the server is marked Down is not always exactly 10 seconds, and can vary depending on how frequently subscribers are logging in. When subscribers are continually and rapidly logging in, the server is marked as Down at 10 seconds. However, if subscribers are logging in less frequently and at a slower pace, then the server is not marked Down until a subsequent subscriber attempts to log in. • All servers cannot be marked as DOWN; instead, the unresponsive servers are marked as UNREACHABLE. For example, if only one RADIUS server is configured and that server is unresponsive, the server status is marked as UNREACHABLE rather than DOWN.	All levels

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• If at least one server has a status of UP, the status of all unresponsive servers is set to DOWN for the remainder of the configured revert-interval setting. • If no server has a status of UP, then the status of the unresponsive servers is set to UNREACHABLE for the remainder of the revert-interval setting or for 30 seconds, whichever is less. • The status of unresponsive servers is returned to UP from DOWN or UNREACHABLE at the end of the revert-interval setting (or the 30-second interval). • If no requests are sent to a server, the server's status is always UP.	
RADIUS servers	Details for specific RADIUS server, identified by IP address.	Detail
Authentication requests	Number of authentication requests received by the authentication server.	Detail
Authentication rollover requests	Number of requests coming into the server as a result of the previous server timing out.	Detail
Authentication retransmissions	Number of retransmissions.	Detail
Accepts	Number of authentication requests accepted by the authentication server.	Detail
Rejects	Number of authentication requests rejected by the authentication server.	Detail
Challenges	Number of authentication requests challenged by the authentication server.	Detail

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Authentication bad authenticators`	Number of responses in which the authenticator is incorrect for the authentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Authentication requests pending`	Number of authentication requests waiting for a response.	Detail
`Authentication request timeouts`	Number of times an authentication request to the server timed out.	Detail
`Authentication unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Authentication packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
`Preauthentication requests`	Number of preauthentication requests received by the preauthentication server.	Detail
`Preauthentication rollover requests`	Number of preauthentication requests coming into the server as a result of the previous server timing out.	Detail
`Preauthentication retransmissions`	Number of retransmissions of preauthentication requests.	Detail
`Preauthentication Accepts`	Number of preauthentication requests accepted by the preauthentication server.	Detail

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Preauthentication Rejects	Number of preauthentication requests rejected by the preauthentication server.	Detail
Preauthentication Challenges	Number of preauthentication requests challenged by the preauthentication server.	Detail
Preauthentication malformed responses	Number of responses to preauthentication requests with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
Preauthentication bad authenticators	Number of responses in which the authenticator is incorrect for the preauthentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
Preauthentication requests pending	Number of preauthentication requests waiting for a response.	Detail
Preauthentication request timeouts	Number of times a preauthentication request to the server timed out.	Detail
Preuthentication unknown responses	Number of unknown responses during the preauthentication phase. The RADIUS response type in the header is invalid or unsupported.	Detail
Preauthentication packets dropped	Number of preauthentication packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
Accounting start requests	Number of accounting start requests received.	Detail
Accounting interim requests	Number of accounting interim requests received.	Detail

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop requests`	Number of accounting stop requests received.	Detail
`Accounting rollover requests`	Number of requests coming into the server as a result of the previous server timing out.	Detail
`Accounting retransmissions`	Number of retransmissions.	Detail
`Accounting start responses`	Number of accounting start responses sent by the server.	Detail
`Accounting interim responses`	Number of accounting interim responses sent by the server.	Detail
`Accounting stop responses`	Number of accounting stop responses sent by the server.	Detail
`Accounting malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Accounting bad authenticators`	Number of responses in which the authenticator is incorrect for the accounting request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Accounting requests pending`	Number of accounting requests waiting for a response.	Detail
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out.	Detail

Table 140: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Accounting packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail

Sample Output

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile1
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
Profile: xyz-profile2
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: DOWN ( 60 seconds )
Profile: xyz-profile11
    Server address: 2001:DB8:0:f101::2
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
```

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile3
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UNREACHABLE
Profile: xyz-profile3
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
      Preauthentication port: 1810
      Status: UNREACHABLE
```

show network-access aaa radius-servers detail

```
user@host> show network-access aaa radius-servers detail
Profile: xyz_profile5
    Server address: 192.168.30.188
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Accounting retry: 5
      Accounting port: 60
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.192
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
```

```
      Accounting retry: 5
      Accounting port: 60
      Status: UP
    Server address: 192.168.30.192
      Authentication port: 1812
      Accounting port: 1813
      Status: UP

RADIUS Servers
  192.168.30.188
    Authentication requests: 7658
    Authentication rollover requests: 0
    Authentication retransmissions: 3600
    Accepts: 6458
    Rejects: 0
    Challenges: 0
    Authentication malformed responses: 0
    Authentication bad authenticators: 0
    Authentication requests pending: 0
    Authentication request timeouts: 4800
    Authentication unknown responses: 0
    Authentication packets dropped: 0
    Preauthentication requests: 7658
    Preauthentication rollover requests: 0
    Preauthentication retransmissions: 3600
    Preauthentication Accepts: 6458
    Preauthentication Rejects: 0
    Preauthentication Challenges: 0
    Preauthentication malformed responses: 0
    Preauthentication bad authenticators: 0
    Preauthentication requests pending: 0
    Preauthentication request timeouts: 4800
    Preauthentication unknown responses: 0
    Preauthentication packets dropped: 0
    Accounting start requests: 1
    Accounting interim requests: 1
    Accounting stop requests: 0
    Accounting rollover requests: 0
    Accounting retransmissions: 0
    Accounting start responses: 1
    Accounting interim responses: 1
    Accounting stop responses: 0
    Accounting malformed responses: 0
```

```
    Accounting bad authenticators: 0
    Accounting requests pending: 0
    Accounting request timeouts: 0
    Accounting unknown responses: 0
    Accounting packets dropped: 0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information | 601

show network-access aaa statistics

IN THIS SECTION

- Syntax | 2958
- Description | 2959
- Options | 2959
- Required Privilege Level | 2959
- Output Fields | 2959
- Sample Output | 2970
- Release Information | 2974

Syntax

```
show network-access aaa statistics
<accounting (detail)>
<address-assignment (client | pool pool-name)>
<dynamic-requests>
```

```
<radius>
<session-limit-per-username>
```

Description

Display AAA accounting, address-assignment, dynamic request statistics, RADIUS settings and statistics, and subscriber session limit statistics.

Options

accounting (detail)	(Optional) Display AAA accounting statistics. The `detail` keyword displays additional accounting information
address-assignment (client | pool *pool-name*)	(Optional) Display AAA address-assignment client and pool statistics.
dynamic-requests	(Optional) Display AAA dynamic requests.
radius	(Optional) Display RADIUS settings and statistics.
session-limit-per-username	Maximum number of sessions allowed for a username per access profile. Use the `brief` option to display only active users with blocked requests. Use the `detail` option to display all active users.

Required Privilege Level

view

Output Fields

Table 141 on page 2960 lists the output fields for the `show network-access aaa statistics` command. Output fields are listed in the approximate order in which they appear.

Table 141: show network-access aaa statistics Output Fields

Field Name	Field Description	Level of Output
`Requests received`	• Number of accounting requests generated by the AAA framework. • Number of dynamic requests received from the external server. Does not include requests sent from backup accounting.	All levels
`Accounting request failures`	Number of accounting requests that failed to be sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting request success`	Number of accounting requests successfully sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Account on requests`	Number of accounting on requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting start requests`	Number of accounting start requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting interim requests`	Number of accounting interim requests sent from a client to a RADIUS accounting server.	`detail`

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop requests`	Number of accounting stop requests sent from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out. This field was named `Timed out requests` in releases before Junos OS Release 16.1. Does not include requests sent from backup accounting.	All levels
`Accounting Response failures`	Number of accounting requests not acknowledged (NAK) by the accounting server. Does not include requests sent from backup accounting.	All levels
`Accounting response success`	Number of accounting requests acknowledged by the accounting server. Does not include requests sent from backup accounting.	All levels
`Account on responses`	Number of accounting on requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting start responses`	Number of accounting start requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting interim responses`	Number of accounting interim requests acknowledged by the RADIUS accounting server.	`detail`

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop responses`	Number of accounting stop requests acknowledged by the RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting rollover requests`	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.	`detail`
`Accounting unknown requests`	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, when the header has invalid or unsupported information).	`detail`
`Accounting radius pending requests`	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.	`detail`
`Accounting malformed responses`	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.	`detail`
`Accounting retransmissions`	Number of accounting requests made by a client to the RADIUS sever that were retransmitted. Does not include requests sent from backup accounting.	`detail`
`Accounting bad authenticators`	Number of accounting responses from a RADIUS accounting server that have an incorrect authenticator (for example, the client and server RADIUS secret do not match).	`detail`

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting packets dropped	Number of accounting responses from a RADIUS accounting server that are dropped by a client.	detail
Accounting backup record creation requests	Number of accounting stop requests from a client to a RADIUS accounting server that were forwarded to be backed up.	detail
Accounting backup replay request success	Number of backup accounting stop requests successfully created by clients after each timeout for replay to a RADIUS accounting server.	detail
Accounting backup request failures	Number of backup accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.	detail
Accounting backup request success	Number of backup accounting requests successfully sent or queued from a client to a RADIUS accounting server.	detail
Accounting backup timeouts	Number of backup accounting requests that timed out after being sent to a RADIUS accounting server.	detail

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting backup in-flight requests`	Number of backup accounting requests that were successfully sent or queued to a RADIUS accounting server for which no response or error has been received yet. Backup requests are replayed only in the following circumstances: • When the request being replayed receives a positive response, the next request can be replayed. • When the request being replayed receives a timeout response, it can be replayed again. Consequently this intermediate timer displays 1 or 0. The value eventually drops to 0 as requests are responded to positively or fail due to error.	`detail`
`Accounting backup responses success`	Number of backup records that were successfully acknowledged with a positive response from a RADIUS accounting server.	`detail`
`Accounting backup radius requests`	Number of backup requests sent to UDP level. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. An observation that the value is increasing is more significant than the exact value of the counter.	`detail`

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting backup radius responses`	Number of responses received at the UDP level for backup requests. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`
`Accounting backup radius timeouts`	Number of backup requests that timed out after being sent to UDP. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`
`Accounting backup radius pending requests`	Number of backup requests sent to a RADIUS accounting server that are waiting for a response from the server. This is an intermediate state counter that eventually drops to zero as requests are responded to or failed due to error.	`detail`
`Accounting backup radius retransmissions`	Sum of backup request retransmissions for each RADIUS accounting server. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting backup malformed responses	Sum of malformed responses received for backup requests sent to each RADIUS accounting server at the UDP level.	detail
Accounting backup bad authenticators	Sum of responses received for backup accounting requests for each RADIUS accounting server where authenticators were mismatched.	detail
Accounting backup responses dropped	Sum of responses for backup accounting requests for each RADIUS accounting server that were dropped due to various sanity checks.	detail
Accounting backup rollover requests	Sum of backup accounting requests rolled over for each RADIUS accounting server.	detail
Accounting backup unknown responses	Sum of unknown responses for backup accounting requests for each RADIUS accounting server.	detail
Client	Client type; for example, DHCP, Mobile IP, PPP.	none specified
Out of Memory	Number of times an address was not given to the client due to memory issues.	none specified
No Matches	Number of times there were no network matches for the pool.	none specified

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pool Name`	Name of the address-assignment pool for this client.	none specified
`Out of Addresses`	Number of times there were no available addresses in the pool.	none specified
`Address total`	Number of addresses in the pool.	none specified
`Addresses in use`	Number of addresses in use.	none specified
`Addresses excluded`	Number of addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.	none specified
`Address Usage (percent)`	Percentage of total addresses in use. This value does not take excluded addresses into account.	none specified
`Pool drain configured`	Configuration state of active drain for the specified local address pool, `yes` or `no`.	none specified
`Pool Usage`	Percentage of allocated addresses in the specified address pool.	none specified
`processed successfully`	Number of dynamic requests processed successfully by the AAA framework.	All levels
`errors during processing`	Number of dynamic requests that resulted in processing errors by the AAA framework.	All levels

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link Name`	Name of the secondary address-assignment pool to which the primary pool is linked.	
`silently dropped`	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.	All levels
`RADIUS Server`	IPv4 or IPv6 address of the RADIUS server to which the router is sending requests.	All levels
`Profile`	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.	All levels
`Configured`	Configured maximum number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded. The range of values is 0 through 2000 outstanding requests. The default value is 1000.	All levels
`Current`	Current number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded.	All levels

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Peak	Highest number of outstanding requests from the router to the RADIUS server for a specific profile at any point in time since the router was started or since the counter was last cleared. **NOTE**: If the value of this field is equal to the value of the Configured field, you may want to increase the value of the Configured field.	All levels
Exceeded	Number of times that the router attempted to send requests to the RADIUS server in excess of the configured maximum value for a specific profile. **NOTE**: If the value of this field is nonzero, you may want to increase the value of the Configured field.	All levels
Username	Username for a subscriber with one or more active sessions for an access profile.	briefdetail
Access-profile	Name of the access profile where the username is active.	briefdetail
Blocked requests	Number of session requests that have been blocked for the username for an access profile. A request is blocked when it exceeds the configured session limit.	briefdetail
Session count	Number of active sessions for the username for an access profile.	briefdetail
Total usernames	Number of active usernames for all access profiles.	none summary

Table 141: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Total usernames exceeding session limit	Number of usernames that have attempted sessions greater than the limit configured for the username.	none summary
Total blocked requests	Number of session requests that have been blocked because the session limit is exceeded.	none summary

Sample Output

show network-access aaa statistics accounting

```
user@host> show network-access aaa statistics accounting
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
```

show network-access aaa statistics accounting detail

```
user@host> show network-access aaa statistics accounting detail
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request failures: 0
  Accounting request success: 5000
    Account on requests: 0
    Accounting start requests: 3000
    Accounting interim requests: 0
    Accounting stop requests: 2000
```

```
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
    Account on responses: 0
    Accounting start responses: 3000
    Accounting interim responses: 0
    Accounting stop responses: 0
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting radius pending requests: 0
  Accounting malformed responses: 0
  Accounting retransmissions: 6000
  Accounting bad authenticators: 0
  Accounting packets dropped: 0

  Accounting backup record creation requests: 3000
  Accounting backup request replay success: 9808
  Accounting backup request failures: 0
  Accounting backup request success: 3006
  Accounting backup timeouts: 6
  Accounting backup in-flight requests: 0
  Accounting backup responses success: 3000
  Accounting backup radius requests: 3006
  Accounting backup radius responses: 3000
  Accounting backup radius timeouts: 99
  Accounting backup radius pending requests: 0
  Accounting backup radius retransmissions: 99
  Accounting backup malformed responses: 0
  Accounting backup bad authenticators: 0
  Accounting backup responses dropped: 0
  Accounting backup rollover requests: 0
  Accounting backup unknown responses: 0
```

show network-access aaa statistics address-assignment client

```
user@host> show network-access aaa statistics address-assignment client
Address-assignment statistics
 Client: jdhcpd
 Out of Memory: 0
 No Matches: 2
```

show network-access aaa statistics address-assignment pool

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 9
 Address total: 47
 Addresses in use: 47
 Address Usage (percent): 100
 Pool drain configured: yes
```

show network-access aaa statistics address-assignment pool (Excluded Addresses)

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 0
 Address total: 24000
 Addresses in use: 12000
 Addresses excluded: 1000
 Address Usage (percent): 50
 Pool drain configured: yes
```

show network-access aaa statistics dynamic-requests

```
user@host> show network-access aaa statistics dynamic-requests
requests received: 0
processed successfully: 0
errors during processing: 0
silently dropped: 0
```

show network-access aaa statistics radius

```
user@host> show network-access aaa statistics radius
Outstanding Requests
RADIUS Server            Profile        Configured   Current   Peak   Exceeded
198.51.100.239           prof1          1000         0         1000   14
                         prof2          500          17        432    0
198.51.100.211           myprof         200          0         200    27
203.0.113.254            pppoe-auth     111          0         1      0
2001:db8:0:f101::2       xyz-profile11  1000         10        135    0
```

show network-access aaa statistics session-limit-per-username (Users with Blocked Requests)

```
user@host> show network-access aaa statistics session-limit-per-username brief
Username           Access-profile     Blocked requests     Session count
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
```

show network-access aaa statistics session-limit-per-username (All Active Users)

```
user@host> show network-access aaa statistics session-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
pqr@example.net    BNG2               0                    1
```

show network-access aaa statistics session-limit-per-username

```
user@host> show network-access aaa statistics on-limit-per-username
Total usernames:  15
Total usernames exceeding session limit: 2
Total blocked requests:  5
```

Release Information

Command introduced in Junos OS Release 9.1.

`address-assignment` option introduced in Junos OS Release 10.0.

`radius` option introduced in Junos OS Release 11.4.

`detail` option introduced in Junos OS Release 13.3.

`session-limit-per-username` option introduced in Junos OS Release 18.4R1 on MX Series routers.

RELATED DOCUMENTATION

- Verifying and Managing Subscriber AAA Information | 601
- Session Options for Subscriber Access | 503
- clear network-access aaa statistics | 2602

show network-access aaa statistics authentication

IN THIS SECTION

- Syntax | 2974
- Description | 2975
- Options | 2975
- Required Privilege Level | 2975
- Output Fields | 2975
- Sample Output | 2977
- Release Information | 2978

Syntax

```
show network-access aaa statistics authentication
<detail>
```

Description

Display AAA authentication statistics.

Options

detail (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 142 on page 2975 lists the output fields for the `show network-access aaa statistics authentication` command. Output fields are listed in the approximate order in which they appear.

Table 142: show network-access aaa statistics authentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of authentication requests received from clients.	All levels
`Accepts`	Number of authentication requests accepted by the authentication server.	All levels
`Rejects`	Number of authentication requests rejected by the authentication server.	All levels
`Challenges`	Number of authentication requests challenged by the authentication server.	All levels
`Timed out requests`	Number of authentication requests that timed out.	All levels
`RADIUS authentication failures`	Number of RADIUS authentication requests that have failed.	Detail

Table 142: show network-access aaa statistics authentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue request deleted	Number of queue requests that have been deleted.	Detail
Malformed reply	Number of malformed replies received from the RADIUS authentication server.	Detail
No server configured	Number of authentication requests that failed because no authentication server is configured.	Detail
Access Profile configuration not found	Number of authentication requests that failed because no access profile is configured.	Detail
Unable to create client record	Number of times that the router is unable to create the client record for the authentication request.	Detail
Unable to create client request	Number of times that the router is unable to create the client request for the authentication request.	Detail
Unable to build authentication request	Number of times that the router is unable to build the authentication request.	Detail
No server found	Number of requests to the authentication server that have timed out; the server is then considered to be down.	Detail
Unable to create handle	Number of authentication requests that have failed because of an internal allocation failure.	Detail
Unable to queue request	Number of times the router was unable to queue the request to the authentication server.	Detail
Invalid credentials	Number of times the router did not have proper authorization to access the authentication server.	Detail

Table 142: show network-access aaa statistics authentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Malformed request	Number of times the router request to the authentication server is malformed.	Detail
License unavailable	Number of times the router did not have a license to access the authentication server.	Detail
Redirect requested	Number of authentication requests that have been redirected based on routing instance.	Detail
Internal failure	Number of internal failures.	Detail
Local authentication failures	Number of times local authentication failed.	Detail
LDAP lookup failures	Number of times the LDAP lookup operation failed.	Detail

Sample Output

show network-access aaa statistics authentication

```
user@host> show network-access aaa statistics authentication
Authentication module statistics
  Requests received: 2118
    Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

show network-access aaa statistics authentication detail

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
```

```
    Accepts: 261
  Rejects: 975
    RADIUS authentication failures: 975
      Queue request deleted: 0
      Malformed reply: 0
      No server configured: 0
      Access Profile configuration not found: 0
      Unable to create client record: 0
      Unable to create client request: 0
      Unable to build authentication request: 0
      No server found: 975
      Unable to create handle: 0
      Unable to queue request: 0
      Invalid credentials: 0
      Malformed request: 0
      License unavailable: 0
      Redirect requested: 0
      Internal failure: 0
    Local authentication failures: 0
    LDAP lookup failures: 0
  Challenges: 0
  Timed out requests: 882
```

Release Information

Command introduced in Junos OS Release 9.1.

Option `detail` introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information | **601**

show network-access aaa statistics pending-accounting-stops

IN THIS SECTION

- Syntax | 2979
- Description | 2979
- Options | 2979
- Required Privilege Level | 2979
- Output Fields | 2979
- Sample Output | 2980
- Release Information | 2980

Syntax

```
show network-access aaa statistics pending-accounting-stops
```

Description

Display the number of pending accounting stop requests.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 143 on page 2980 lists the output field for the `show network-access aaa statistics pending-accounting-stops` command.

Table 143: show network-access aaa statistics pending-accounting-stops Output Fields

Field Name	Field Description
`Pending accounting stops`	Total number of accounting stop messages queued.

Sample Output

show network-access aaa statistics pending-accounting-stops

```
user@host> show network-access aaa statistics pending-accounting-stops
Pending accounting stops: 10,000
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- request network-access aaa replay pending-accounting-stops | 2641
- show accounting pending-accounting-stops | 2669

show network-access aaa statistics preauthentication

IN THIS SECTION

- Syntax | 2981
- Description | 2981
- Options | 2981
- Required Privilege Level | 2981
- Output Fields | 2981

Sample Output | 2982
Release Information | 2982

Syntax

```
show network-access aaa statistics preauthentication
```

Description

Display AAA preauthentication statistics.

Options

detail (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 144 on page 2981 lists the output fields for the `show network-access aaa statistics preauthentication` command. Output fields are listed in the approximate order in which they appear.

Table 144: show network-access aaa statistics preauthentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of preauthentication requests received from clients.	All levels
`Multistack requests`	Number of preauthentication requests for dual-stack subscribers.	All levels

Table 144: show network-access aaa statistics preauthentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accepts	Number of preauthentication requests accepted by the preauthentication server.	All levels
Rejects	Number of preauthentication requests rejected by the preauthentication server.	All levels
Challenges	Number of preauthentication requests challenged by the preauthentication server.	All levels
Timed out requests	Number of preauthentication requests that timed out.	All levels

Sample Output

show network-access aaa statistics preauthentication

```
user@host> show network-access aaa statistics preauthentication
Preauthentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

RADIUS Logical Line Identifier (LLID) Overview | 544

Configuring Logical Line Identification (LLID) Preauthentication | 546

show network-access aaa statistics re-authentication

IN THIS SECTION

- Syntax | 2983
- Description | 2983
- Required Privilege Level | 2983
- Output Fields | 2983
- show network-access aaa statistics re-authentication | 2984
- Release Information | 2984

Syntax

```
show network-access aaa statistics re-authentication
```

Description

Display statistics for RADIUS re-authentication, and starting in Junos OS Release 18.2R1, for local re-authentication.

Required Privilege Level

view

Output Fields

Table 145 on page 2984 lists the output fields for the show network-access aaa statistics re-authentication command. Output fields are listed in the approximate order in which they appear.

Table 145: show network-access aaa statistics re-authentication Output Fields

Field Name	Field Description
`Re-authentication statistics`	Displays re-authentication statistics.
`Requests received`	Total number of re-authentication requests that the device received from clients.
`Accepts`	Total number of accepted re-authentications.
`Challenges`	Total number of re-authentication challenges.
`Internal errors`	Total number of re-authentication internal errors.
`Rejects`	Total number of re-authentications rejected.
`Timed out requests`	Total number of re-authentication accounting timeouts.

show network-access aaa statistics re-authentication

command-name

```
user@host> show network-access aaa statistics re-authentication
Re-authentication statistics
  Requests received: 0
  Accepts: 0
  Rejects: 0
  Challenges: 0
  Timed out requests: 0
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

reauthenticate (DHCP Local Server) | 2267

show network-access aaa subscribers

IN THIS SECTION

- Syntax | 2985
- Description | 2985
- Options | 2985
- Required Privilege Level | 2986
- Output Fields | 2986
- Sample Output | 2989
- Release Information | 2992

Syntax

```
show network-access aaa subscribers
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<statistics>
<username>
<session-id session-id-number detail>
```

Description

Display subscriber-specific AAA statistics.

Options

logical-system *logical-system-name*	(Optional) List subscribers in the specific logical system.

routing-instance *routing-instance-name*	(Optional) List subscribers for the specific routing instance. If you do not specify a routing instance name, the default routing instance is assumed.
statistics	(Optional) Display statistics for the subscriber events.
username	(Optional) Display information for the specified subscriber.
session-id *session-id-number* **detail**	(Optional) Display information for the specified session ID.

Required Privilege Level

view

Output Fields

Table 146 on page 2986 lists the output fields for the `show network-access aaa subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 146: show network-access aaa subscribers Output Fields

Field Name	Field Description
`Challenge requests`	Number of authentication requests challenged by the authentication server for this subscriber.
`Challenge responses`	Number of challenge responses sent by the subscriber to the authentication server.
`START sent successfully`	Number of accounting start requests generated by the AAA framework for this subscriber.
`START send failures`	Number of accounting start requests that failed to make it to the accounting server for this subscriber.
`START ack received`	Number of accounting start requests acknowledged by the accounting server for this subscriber.
`INTERIM sent successfully`	Number of accounting interim requests generated by the AAA framework for this subscriber.

Table 146: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
INTERIM send failures	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
INTERIM ack received	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
Requests received	Number of reauthentication requests received by the authentication server.
Successful responses	Number of successful reauthentication requests granted by the authentication server.
Aborts handled	Number of reauthentication requests terminated by the authentication server.
Service name	Name of the subscriber service.
Creation requests	Number of requests to create the service.
Deletion requests	Number of requests to delete the service.
Request timeouts	Number of times the service request was timed out.
Client type	Type of client; for example, DHCP, Mobile IP, PPP.
Session-ID	ID of the subscriber session.
Session uptime	How long the session has been up, in *HH:MM:SS.*
Accounting	Status of accounting, and type of accounting if accounting is on.
Stripped username	Username of the subscriber session.

Table 146: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
AAA Logical system/ Routing instance	AAA framework for the subscriber of logical system or routing instance.
Target Logical system/Routing instance	Target framework for the subscriber of logical system or routing instance.
Access-profile	Profile of the subscriber.
Accounting Session ID	ID of the subscriber session for accounting.
Multi Accounting Session ID	ID of the subscriber session for multiple accouting.
IP Address	IPv4 address of the subscriber.
IPv6 Address	IPv6 address of the subscriber.
IPv6 Prefix	IPv6 prefix of the subscriber.
Authentication State	State of subscriber session authentication.
Accounting State	State of subscriber session accounting.
Provisioning Type	Type of subscriber provisioning.

Sample Output

show network-access aaa subscribers logical-system

```
user@host> show network-access aaa subscribers logical-system
Username            Client type    Logical system/Routing instance
user61@example.net    ppp            default
00010e020304.1231   dhcp           isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp           default:isp-gtown-r3-00
0020df980102.2334   dhcp           isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers logical-system routing-instance

```
user@host> show network-access aaa subscribers logical-system isp-bos-metro-16 routing-instance
isp-cmbrg-12-32
Username            Client type   Logical system/Routing instance
00010e020304.1231   dhcp          isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp          default:isp-gtown-r3-00
0020df980102.2334   dhcp          isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers statistics username

```
user@host> show network-access aaa subscribers statistics username 00010e020304.1231
Authentication statistics
   Challenge requests: 0
   Challenge responses: 0
 Accounting statistics
   START sent successfully: 1
   START send failures: 0
   START ack received: 1
   INTERIM sent successfully: 0
   INTERIM send failures: 0
   INTERIM ack received: 0
 Re-authentication statistics
   Requests received: 0
   Sucessfull responses: 0
   Aborts handled: 0
 Service statistics
```

```
   Service name: filter-serv
   Creation requests: 1
   Deletion requests: 0
   Request timeouts: 0
   Service name: filter-serv2
   Creation requests: 144
   Deletion requests: 0
   Request timeouts: 144
```

show network-access aaa subscribers username

```
user@host> show network-access aaa subscribers username user80@example.net
Logical system/Routing instance   Client type  Session-ID  Session uptime  Accounting
isp-bos-metro-16:isp-cmbrg-12     dhcp         7           01:12:56        on/volume
Service name       Service type      Quota          Accounting
I-Cast             volume            1200 Mbps      on/volume+time
Voip                                                on/volume
GamingBurst        time              6000 secs      on/volume
```

show network-access aaa subscribers session-id 26 detail

The following command output is seen when only an IPv4 client is associated with the session:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

The following command output is seen when IPv6 client logs in (after IPv4 association) and is associated with the same session ID:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
IPv6 Address: 2001:db8:0:8003::2
IPv6 Prefix: 2001:db8:ffff:0:4::/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show network-access aaa subscribers (Tenant systems)

```
user@host:TSYS1> show network-access aaa subscribers
Username                Logical system/Routing instance   Client type    Session-ID
userX                   default:TSYS1-ri                    xauth          1
```

show network-access aaa subscribers (Tenant systems)

```
user@host> show network-access aaa subscribers session-id 1 detail
Type: xauth
Username: userX
Stripped username: userX
AAA Logical system/Routing instance: default:TSYS1-ri
Target Logical system/Routing instance: default:TSYS1-ri
Access-profile: ap1
+    Tenant: TSYS1
Session ID: 1
Multi Accounting Session ID: 0
IP Address: 192.0.2.0
Authentication State: AuthStateActive
```

```
Accounting State: Acc-Init
Converted to time accounting: no
Provisioning Type: None

user@host> show network-access aaa subscribers session-id 2 detail
Type: xauth
Username: userY
Stripped username: userY
AAA Logical system/Routing instance: default:TSYS2-ri
Target Logical system/Routing instance: default:TSYS2-ri
Access-profile: ap1
+    Tenant: TSYS2
Session ID: 2
Multi Accounting Session ID: 0
IP Address: 192.0.2.1
Authentication State: AuthStateActive
Accounting State: Acc-Init
Converted to time accounting: no
Provisioning Type: None
```

Release Information

Command introduced in Junos OS Release 9.1.

Command updated with `session-id` *`session-id-number`* `detail` in Junos OS Release 17.3.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information | 601

show network-access aaa subscribers session-id

IN THIS SECTION

- Syntax | 2993
- Description | 2993

- Options | 2993
- Required Privilege Level | 2993
- Output Fields | 2993
- Sample Output | 2999
- Release Information | 3003

Syntax

```
show network-access aaa subscribers session-id session-id
<brief | detail>
```

Description

Display information about the specified subscriber session.

Options

session-id	ID of the subscriber session.
brief | detail	(Optional) Display the specified level of information.

Required Privilege Level

view

Output Fields

Table 147 on page 2994 lists the output fields for the `show network-access aaa subscribers session-id` command. Output fields are listed in the approximate order in which they appear.

Table 147: show network-access aaa subscribers session-id Output Fields

Field Name	Field Description	Level of Output
Type and Client type	Type of client.	All levels
Accounting	Status of the accounting configuration for the service, on or off, and the type of accounting, time, volume+time, or flat-file. The time and volume+time types are configured in RADIUS Service-Statistics VSA [26-69].	brief none
Service type	Type of accounting: volume, time, volume+time, or na.	brief
Quota	Quota for service: volume (in Mbps) or time (seconds).	brief
Username	Name of the user logged in to the session.	detail
Stripped username	Username after the domain has been removed.	detail
Logical system/ Routing instance and AAA Logical system/ Routing instance	Name of the routing instance, logical system name, or both used for the session.	All levels
Target Logical system/Routing instance	Logical system/routing instance to which the session is mapped.	detail
Access-profile	Access profile used for AAA services for the session.	detail
Session ID	ID of the subscriber session. The session ID value displayed under Service name is the service session ID.	detail

Table 147: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail
`Multi Accounting Session ID`	Bundle ID for MLPPP sessions. Acct-Multi-Session-Id (RADIUS attribute 50) uses the value of the session database bundle session ID to enable RADIUS to link together multiple related sessions. The value of this field is zero when no MLPPP sessions exist.	detail
`IP Address`	IP address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`IPv6 Address`	IPv6 address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`IPv6 Prefix`	IPv6 prefix of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`Authentication State`	State of the subscriber authentication session: `AuthInit`, `AuthStart`, `AuthChallenge`, `AuthRedirect`, `AuthClntRespWait`, `AuthAcctVolStatsAckWait`, `AuthAcctStopAckWait`, `AuthServCreateRespWait`, `AuthLogoutStart`, `AuthStateActive`, `AuthClntLogoutRespWait`, `AuthProfileUpdateWait`, `AuthProvisionRespWait`, `AuthProvisionServiceCreationWait`	detail
`Gx-Plus Provisioning State`	State of Gx-Plus provisioning: • ignored—Subscriber has no IPv4 address or NAS-Port-ID. • in-progress—Provisioning is in progress. • logout—Subscriber logout is in progress. • logout-done—Logout response has been received. • response-received—Provisioning response has been received.	detail

Table 147: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pcrf Provisioning State`	State of PCRF provisioning: • `active`—PCRF provisioning is active. • `ignored`—Subscriber has no IPv4 address or NAS-Port-ID. • `in-progress`—Provisioning is in progress. • `logout`—Subscriber logout is in progress. • `logout-done`—Logout response has been received. • `response-received`—Provisioning response has been received.	detail
`Pcrf Subscription-Id-Type`	Type of subscriber for a PCRF partition. You can define your own or use a predefined value: `0 (END_USER_E164), 1 (END_USER_IMSI), 2 (END_USER_SIP_URI), 3 (END_USER_NAI), 4 (END_USER_PRIVATE)`.	detail
`Pcrf Subscription-Id-Data`	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for a PCRF partition; for example, demux0	detail
`Ocs Subscription-Id-Type`	Type of subscriber for an OCS partition. You can define your own or use a predefined value: `0 (END_USER_E164), 1 (END_USER_IMSI), 2 (END_USER_SIP_URI), 3 (END_USER_NAI), 4 (END_USER_PRIVATE)`.	detail
`Ocs Subscription-Id-Data`	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for an OCS partition; for example: test-sid	detail
`Ocs Interrogation State`	State of the OCS interrogation: `first, intermediate, final`.	detail
`Ocs Data State`	State of the OCS data: `none`	detail

Table 147: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting State`	State of the subscriber accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	detail
`Provisioning-type`	Provisioning type for this session: • gx-plus—Subscriber service uses Gx-Plus provisioning. • jsrc—Subscriber service uses JSRC provisioning. • none—Provisioning is not enabled.	detail
`Service name`	Name of the attached service or policy. • For RADIUS-activated and CLI-activated services, displays the full activation string for the service. If the activation string includes service parameters, then both the service name and service parameters are displayed. • For JSRC-activated policies—displays the policy name.	All levels
`Service State`	State of the service provided in the subscriber session.	detail
`Service Family`	Network family of the service provided in the subscriber session.	detail
`Service Activation Source`	Source used to activate the service.	detail
`Session uptime`	How long the session has been up, in *HH:MM:SS*.	All levels
`Service CC-Service-Identifier`	Data identification element of the 3GPP Diameter credit control service charging system that uniquely defines the `CC-Service-Context`.	detail

Table 147: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Service Rating-Group`	Value associated with a charging rule and part of the accounting data stream for the PCRF.	detail
`Ocs Control`	Whether OCS controls the service: `yes` or `no`.	detail
`Accounting status`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time`, `volume+time`, or `flat-file`. The `time` and `volume+time` types are configured in RADIUS Service-Statistics VSA [26-69].	detail
`Service accounting session ID`	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail
`Service accounting state`	State of the service accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	detail
`Accounting interim interval`	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail
`Pcrf session-stamp`	Value that consists of the UTC time when the router creates the CCR-GX-I. The router appends the session stamp to the session ID when a CCR packet is sent to the PCRF. You configure the `use-session-stamp` option at the `[edit access pcrf partition` *`partition-name`* hierarchy level. If you do not configure this option, the field displays a value of zero. For local reinitialization, you must configure the `use-session-stamp` option.	detail

Sample Output

show network-access aaa subscribers session-id brief

```
user@host> show network-access aaa subscribers session-id 6 brief
Logical system/Routing instance   Client type   Session uptime    Accounting
default:default                   dhcp          00:01:29          on/time

Service name               Service type   Quota           Accounting
filter-service             -na-           -na-            off
filter-service-2           volume+time    77.00MB/120secs off
1337994190863204450        -na-           -na-            off
```

show network-access aaa subscribers session-id 2 (Flat File Accounting)

```
user@host> show network-access aaa subscribers session-id 2
Logical system/Routing instance   Client type    Session-ID     Session uptime    Accounting
default:default                   dhcp           2              00:00:48          on/volume+time

Service name               Service type   Quota           Accounting
filter-service             -na-           -na-            on/flat-file
```

show network-access aaa subscribers session-id detail

```
user@host> show network-access aaa subscribers session-id 5 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: 5
Accounting Session ID: jnpr ge-1/0/0.101:1
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Pcrf Provisioning State: active
```

```
Pcrf Subscription-Id-Type: 4
Pcrf Subscription-Id-Data: demux0
Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: test-sid
Ocs Interrogation State: intermediate
Ocs Data State: none
Gx-Plus Provisioning State: response-received
Accounting State: Acc-Interim-Sent
Provisioning-type: jsrc
Service name: filter-service-1
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN
 Session ID: 7
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 777
 Service Rating Group: 10
 Ocs Control: yes
Service name: filter-service-2
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN Session ID: 8
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 778
 Service Rating Group: 11
 Ocs Control: no
Accounting status: on/volume+time
 Service accounting session ID: 1:2-1322506006
 Service accounting state: Acc-Interim-Sent
 Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Service with Multiple Instances)

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: user-test-fms2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
```

```
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius at Reauth
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Single Session Dual Stack with active V4 and V6 subscribers)

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: user-test-25
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.10.0.6
IPv6 Address: 00:00:5E:00:53:02
```

```
IPv6 Prefix: 00:00:5E:00:53:00/64
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show network-access aaa subscribers session-id detail (PCRF with Session-Stamp)

```
user@host> show network-access aaa subscribers session-id 23 detail
Type: dhcp
Username: user45-28-abc@example.net
Stripped username: user45-28-abc
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: test-profile-abc-user
Session ID: 23
Accounting Session ID: 23
Multi Accounting Session ID: 0
IP Address: 198.51.100.5
Authentication State: AuthStateActive
Pcrf session-stamp: 1557788595
Pcrf Provisioning State: active
Pcrf Subscription-Id-Type: 4
Pcrf Subscription-Id-Data: user45-28-abc@example.net
Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: 987654321
Ocs Interrogation State: idle
Ocs Data State: none
Accounting State: Acc-Start-Sent
Provisioning Type: PCRF-LOGIN
Service name: test-filters(filter-123-xyz-in,filter-123-xyz-out)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: PCRF-LOGIN
  Session ID: 24
  Session uptime: 00:00:39
  Service CC-Service-Identifier: 12345
  Service Rating Group: 3300
  Ocs Control: yes
  Service session type: Service-Profile
Service name: test-cos(abc-video-100M)
```

```
  Service State: SvcActive
  Service Activation Source: PCRF-LOGIN
  Session ID: 25
  Session uptime: 00:00:39
  Service session type: Service-Profile
Service name: test-multi
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: PCRF-LOGIN
  Session ID: 26
  Session uptime: 00:00:39
  Service session type: Service-Profile
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- Verifying and Managing Subscriber AAA Information | **601**
- *Local and Remote Service Activation and Deactivation Using the CLI*
- *Deactivating a Single Instance of a Subscriber Service*
- *Deactivating All Instances of a Subscriber Service*
- *Verifying Subscriber Services with Multiple Instances*
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

show network-access aaa terminate-code

IN THIS SECTION

- Syntax | **3004**
- Description | **3004**

- Options | 3004
- Required Privilege Level | 3005
- Output Fields | 3005
- Sample Output | 3006
- Release Information | 3009

Syntax

```
show network-access aaa terminate-code
<brief | detail | summary>
<reverse>
<(aaa | dhcp | l2tp | ppp)>
```

Description

Display the count for termination cause types and the current mapping between session termination cause types and code values.

Options

none	Display all mappings.
brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default and includes base count information about mappings. The `brief` output displays mappings with non-zero usage count and custom mappings. The `detail` output displays all mappings.
aaa	(Optional) Limit display to AAA mappings only.
dhcp	(Optional) Limit display to DHCP mappings only.
l2tp	(Optional) Limit display to L2TP mappings only.
ppp	(Optional) Limit display to PPP mappings only.
reverse	(Optional) Display mapping of the code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) to the termination cause type.

vlan (Optional) Limit display to VLAN mappings only.

Required Privilege Level

view

Output Fields

Table 148 on page 3005 lists the output fields for the `show network-access aaa terminate-code` command. Output fields are listed in the approximate order in which they appear.

Table 148: show network-access aaa terminate-code Output Fields

Field Name	Field Description	Level of Output
RADIUS	RFC-defined code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) or a nonstandard, customized value that you configure with the `terminate-code aaa` statement at the `[edit access]` hierarchy level.	`brief detail` None (with `reverse` option)
Custom	Whether or not the termination cause is a customized mapping or the default mapping.	All levels
Mapping-Count	Number of mappings that occurred for a specific terminate cause type or category (standard or summary output) or per termination cause (reverse output).	`summary` None
Usage-Count	Number of times the terminate code mapping was used.	All levels
Type	Termination cause type—null, aaa, dhcp, l2tp, ppp, or vlan. **NOTE**: The null termination cause type indicates that no termination reason was provided by the subscriber and the RADIUS Acct-Terminate-Cause attribute (49) was not included in the Acct-Stop request	All levels

Table 148: show network-access aaa terminate-code Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Code	Specific termination cause.	`brief detail`

Sample Output

show network-access aaa terminate-code

```
user@host> show network-access aaa terminate-code
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code reverse

```
user@host> show network-access aaa terminate-code reverse
Terminate-code:
  RADIUS    Custom Mapping-Count Usage-Count Type
  0         no     1             0           null
  1         no     1             0           aaa
  1         no     1             0           dhcp
  1         no     5             0           l2tp
  1         no     8             0           ppp
  1         no     2             10          vlan
  2         no     1             0           dhcp
  2         no     3             0           ppp
  2         no     2             0           vlan
  4         no     1             0           aaa
  4         no     1             0           dhcp
  4         no     1             0           l2tp
  4         no     1             0           ppp
```

```
  5         no     2             0           aaa
  5         no     1             0           l2tp
  5         no     1             0           ppp
  6         no     2             0           aaa
  6         no     13            0           l2tp
  6         no     3             0           ppp
  6         no     3             0           vlan
  8         no     3             0           l2tp
  8         no     5             0           ppp
  9         no     13            0           l2tp
  9         no     12            0           ppp
  9         no     4             0           vlan
  10        no     4             0           aaa
  10        no     1             0           dhcp
  10        no     128           0           l2tp
  10        no     171           0           ppp
  15        no     1             0           dhcp
  15        no     190           0           l2tp
  15        no     1             0           vlan
  16        no     1             0           vlan
  17        no     2             0           aaa
  17        no     10            0           l2tp
  17        no     6             0           ppp
```

show network-access aaa terminate-code dhcp

```
user@host> show network-access aaa terminate-code dhcp
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     5             0           dhcp
```

show network-access aaa terminate-code detail

```
user@host> show network-access aaa terminate-code aaa detail

Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  17         no     0           aaa  deny-server-request-timeout
```

```
  6          no     0           aaa  service-shutdown-network-logout
  10         no     0           aaa  service-shutdown-remote-reset
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  5          no     0           aaa  service-shutdown-time-limit
  10         no     0           aaa  service-shutdown-volume-limit
  6          no     13          aaa  shutdown-administrative-reset
  4          no     0           aaa  shutdown-idle-timeout
  10         no     0           aaa  shutdown-remote-reset
  5          no     0           aaa  shutdown-session-timeout
```

show network-access aaa terminate-code brief

```
user@host> show network-access aaa terminate-code brief
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  6          no     13          aaa  shutdown-administrative-reset
  15         no     7           dhcp nak
  10         no     1           l2tp session-receive-cdn-avp-missing-secret
  10         no     1           ppp  bundle-fail-create
  1          no     1           ppp  lcp-peer-terminate-term-req
  10         no     1           ppp  lcp-tunnel-disconnected
  1          no     10          vlan out-of-band-ancp-port-down
```

show network-access aaa terminate-code summary

```
user@host> show network-access aaa terminate-code summary

Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code vlan

```
user@host> show network-access aaa terminate-code vlan
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     13            0           vlan
```

show network-access aaa terminate-code vlan detail

```
user@host> show network-access aaa terminate-code vlan detail
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  6          no     0           vlan admin-logout
  16         no     0           vlan admin-reconnect
  9          no     0           vlan other
  2          no     0           vlan out-of-band-access-interface-down
  6          no     0           vlan out-of-band-admin-access-interface-down
  6          no     0           vlan out-of-band-admin-core-interface-down
  1          no     0           vlan out-of-band-ancp-port-down
  1          no     0           vlan out-of-band-ancp-port-vlan-id-change
  2          no     0           vlan out-of-band-core-interface-down
  15         no     0           vlan out-of-band-l2-wholesale-no-free-vlans
  9          no     0           vlan profile-request-error
  9          no     0           vlan sdb-error
  9          no     0           vlan subscriber-activate-error
```

Release Information

Command introduced in Junos OS Release 11.4.

`vlan` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Understanding Session Termination Causes and RADIUS Termination Cause Codes | 603
- Mapping Session Termination Causes to Custom Termination Cause Codes | 606
- AAA Termination Causes and Code Values | 608
- DHCP Termination Causes and Code Values | 610

L2TP Termination Causes and Code Values | 611
PPP Termination Causes and Code Values | 638
VLAN Termination Causes and Code Values | 651

show network-access address-assignment pool

IN THIS SECTION

- Syntax | 3010
- Description | 3010
- Options | 3010
- Required Privilege Level | 3011
- Output Fields | 3011
- Sample Output | 3012
- Release Information | 3012

Syntax

```
show network-access address-assignment pool pool-name
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display state information for each address-assignment pool.

Options

none	Display information about clients that have obtained addresses from the address-assignment pool.
pool *pool-name*	Display information about the specified address-assignment pool.

logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance.

Required Privilege Level

view and system

Output Fields

Table 149 on page 3011 lists the output fields for the `show network-access address-assignment pool` command. Output fields are listed in the approximate order in which they appear.

Table 149: show network-access address-assignment pool Output Fields

Field Name	Field Description
`IP address/prefix`	IP address of the client.
`Hardware address`	MAC address of the client. Displays `NA` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.
Host/User	Hostname or username of the client. Displays `EXCLUDED` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.
`Type`	Type of client. Displays `unknown` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.

Sample Output

show network-access address-assignment pool

```
user@host> show network-access address-assignment pool sunnywest logical-system ls1 routing-
instance routinst2
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
```

Release Information

Command introduced in Junos OS Release 9.0.

show network-access domain-map

IN THIS SECTION

- Syntax | 3012
- Description | 3013
- Options | 3013
- Required Privilege Level | 3013
- Output Fields | 3013
- Sample Output | 3014
- Release Information | 3014

Syntax

```
show network-access domain-map
<statistics>
```

Description

Display domain map information.

Options

statistics (Optional) Display domain map statistics.

Required Privilege Level

view

Output Fields

Table 150 on page 3013 lists the output fields for the `show network-access domain-map statistics` command. Output fields are listed in the approximate order in which they appear.

Table 150: show network-access domain-map Output Fields

Field Name	Field Description
`Matched domains`	Number of usernames with domain names that are matched.
`Unmatched domains`	Number of usernames with domain names that are not matched.
`Missing domain names`	Number of usernames without a domain name.
`Stripped username`	Number of usernames from which the domain name has been stripped.
`Default used`	Number of times the default domain map is used.

Sample Output

show network-access domain-map statistics

```
user@host> show network-access domain-map statistics
General domain mapping statistics
  Matched domains: 7
  Unmatched domains: 1
  Missing domain names: 0
  Stripped username: 7
Domain statistics for domain-name: default
  Default used: 1
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Verifying Domain Maps | 673

show network-access gx-plus

IN THIS SECTION

- Syntax | 3015
- Description | 3015
- Options | 3015
- Required Privilege Level | 3015
- Output Fields | 3015
- Sample Output | 3018
- Release Information | 3031

Syntax

```
show network-access gx-plus
<state | statistics | sync-state>
```

Description

Display Gx-Plus provisioning state, synchronization state, and statistics information.

Options

state	(Optional) Display Gx-Plus provisioning state.
statistics	(Optional) Display Gx-Plus statistics.
sync-state	(Optional) Display Gx-Plus synchronization state.

Required Privilege Level

view

Output Fields

Table 151 on page 3016 lists the output fields for the `show network-access gx-plus` command. Output fields are listed in the approximate order in which they appear.

Table 151: show network-access gx-plus Output Fields

Field Name	Field Description
`Gx-plus state`	State of the Gx-Plus application, including the following: • `Engine created` • `Partition configured` • `Configuration active` • `Diameter interface configured` • `Total number of entries`—Number of provisioned, pending, and terminating subscribers. • `Number of pending entries`—Number of pending subscribers. • `Number of pending logouts`—Number of subscribers logging out.
`Sync-Event`	Type of synchronization event.
`Timeout`	Number of times notification sent without response.
`Gx-plus general counters`	Number and state of general events.
`Gx-plus sync-event counters`	Number and state of synchronization events.
`Gx-plus sync-event diameter-events`	Events or errors that occur in jdiameterd on behalf of Gx-Plus.
`Gx-plus upstream diameter-events`	Errors encountered when the Gx-Plus component in authd is processing outbound requests.
`Gx-plus upstream counters`	Events or errors that occur with the Gx-Plus component of authd.
`Gx-plus downstream counters`	Events for requests that originated from jdiameterd and are intended for the Diameter layer on a PCRF server.

Table 151: show network-access gx-plus Output Fields *(Continued)*

Field Name	Field Description
cold-boot	Event that takes place after a cold boot (system reboot).
warm-boot	Event that takes place after an authd or jdiameterd process restart.
d-req	Event that takes place after a discovery request from the router.
ayt	Event that takes place after an “Are You There” check. The check is triggered when the connection breaks between Diameter peers. It is the same as an AWD, except it is sent from the router side rather than the server side.
awd	Event that takes place after an application watchdog flag. The AWD is triggered when the connections breaks between Diameter peers. It is the same as an AYT, except it is sent from the server side rather than the router side.
prov	Events that take place as a result of the CCR-I provisioning request.
logout	Events that take place as a result of the CCR-T logout request.
ruleReport	Events that take place as a result of the CCR-U rule report request.
threshold	Events that take place as a result of the CCR-U threshold request.
general	Events that cannot be attributed to a specific type of request.
audit	Events that take place as a result requests for an audit of a specific subscriber session; includes information about state and services.
d-exact	Events that take place as a result of discovery exchange for Gx-Plus. The request originates from the PCRF server for a specific subscriber.

Table 151: show network-access gx-plus Output Fields *(Continued)*

Field Name	Field Description
d-bulk	Events that take place as a result of discovery requests for all active Gx-Plus subscribers. The request originates on the PCRF server.
d-done	Events that take place when the discovery exchange is completed.

Sample Output

show network-access gx-plus state

```
user@host> show network-access gx-plus state
Gx-plus state:
  Engine created                : yes
  Partition configured          : yes
  Configuration active          : yes
  Diameter interface configured : yes
  Total number of entries       : 0
  Number of pending entries     : 0
  Number of pending logouts     : 0
```

show network-access gx-plus statistics

```
user@host> show network-access gx-plus statistics
Gx-plus general counters:
  Counter                                                 Value
  engine created                                          1
  initial config: active                                  1
  recovery: process restart                               1
  diameter-app initial config: success                    1

Gx-plus sync-event counters:
  Category      Counter                                   Value
  warm-boot     activated                                 1
  warm-boot     posted                                    1
  warm-boot     response                                  1
```

```
  awd           posted                                    12
  awd           response                                  12
```

show network-access gx-plus sync-state

```
user@host> show network-access gx-plus sync-state
Gx-plus sync-events:
  Sync-Event                Timeout
  cold-boot                 6100
```

show network-access gx-plus statistics extensive (Extensive)

```
user@host> show network-access gx-plus statistics extensive

Gx-plus general counters:
  Counter                                                 Value
  engine created                                          0
  engine created after recovery                           0
  engine destruction started                              0
  engine destroyed                                        0
  diameter-app config: failed                             0
  initial config: active                                  0
  initial config: inactive                                0
  config changed: inactive to active                      0
  config changed: active to inactive                      0
  config changed: active to active                        0
  config changed: inactive to inactive                    0
  recovery: cold-boot                                     0
  recovery: re-failover                                   0
  recovery: process restart                               0
  diameter-app initial config: failure                    0
  diameter-app initial config: success                    0
  diameter-app config retry: failure                      0
  diameter-app config retry: success                      0
  response when no active config                          0
  bad response                                            0
  response: unknown session                               0
  request when on no origins                              0
  request when no config: response queue full             0
  request when no config: no memory                       0
```

```
  request when no config: bad request                    0
  request when no config: response build failure         0
  request when no config: response sent                  0
  request when no config: response post failure          0
  bad request                                            0
  diameter tx                                            124986
  diameter rx                                            7231
  diameter event                                         1253
  retrieving pending provisioning request                0
  retrieving pending rule-report                         0
  retrieving pending rule-report                         0
  missed rule report request                             0
  tick                                                   1213965

Gx-plus sync-event counters:
  Sync-Event    Counter                                  Value
  cold-boot     activated                                0
  cold-boot     re-activated                             0
  cold-boot     duplicate activation                     0
  cold-boot     activated to queue                       0
  cold-boot     message build failure                    0
  cold-boot     posted                                   0
  cold-boot     post failure                             0
  cold-boot     response                                 0
  warm-boot     activated                                0
  warm-boot     re-activated                             0
  warm-boot     duplicate activation                     0
  warm-boot     activated to queue                       0
  warm-boot     message build failure                    0
  warm-boot     posted                                   0
  warm-boot     post failure                             0
  warm-boot     response                                 0
  d-req         activated                                0
  d-req         re-activated                             0
  d-req         duplicate activation                     0
  d-req         activated to queue                       0
  d-req         message build failure                    0
  d-req         posted                                   0
  d-req         post failure                             0
  d-req         response                                 0
  ayt           activated                                5
  ayt           re-activated                             0
  ayt           duplicate activation                     37
```

```
  ayt           activated to queue                        0
  ayt           message build failure                     0
  ayt           posted                                    0
  ayt           post failure                              0
  ayt           response                                  0
  awd           activated                                 0
  awd           re-activated                              0
  awd           duplicate activation                      0
  awd           activated to queue                        0
  awd           message build failure                     0
  awd           posted                                    0
  awd           post failure                              0
  awd           response                                  0

Gx-plus sync-event diameter-events:
  Category      Counter                                   Value
  cold-boot     bad flags                                 0
  cold-boot     bad fixed destination                     0
  cold-boot     bad routed destination                    0
  cold-boot     tx is over limit                          0
  cold-boot     bad end-to-end id                         0
  cold-boot     no peer for tx                            0
  cold-boot     peer down while waiting for answer        0
  cold-boot     timeout while waiting for answer          0
  cold-boot     tx timeout                                0
  cold-boot     tx try limit                              0
  cold-boot     tx failure                                0
  cold-boot     discarded                                 0
  cold-boot     received answer is over limit             0
  cold-boot     tx failure: no memory                     0
  cold-boot     base-app-tx-timeout                       0
  cold-boot     base-app-rx-timeout                       0
  cold-boot     base-app-tx-discard                       0
  cold-boot     base-app-rx-discard                       0
  warm-boot     bad flags                                 0
  warm-boot     bad fixed destination                     0
  warm-boot     bad routed destination                    0
  warm-boot     tx is over limit                          0
  warm-boot     bad end-to-end id                         0
  warm-boot     no peer for tx                            0
  warm-boot     peer down while waiting for answer        0
  warm-boot     timeout while waiting for answer          0
  warm-boot     tx timeout                                0
```

```
warm-boot     tx try limit                              0
warm-boot     tx failure                                0
warm-boot     discarded                                 0
warm-boot     received answer is over limit             0
warm-boot     tx failure: no memory                     0
warm-boot     base-app-tx-timeout                       0
warm-boot     base-app-rx-timeout                       0
warm-boot     base-app-tx-discard                       0
warm-boot     base-app-rx-discard                       0
d-req         bad flags                                 0
d-req         bad fixed destination                     0
d-req         bad routed destination                    0
d-req         tx is over limit                          0
d-req         bad end-to-end id                         0
d-req         no peer for tx                            0
d-req         peer down while waiting for answer        0
d-req         timeout while waiting for answer          0
d-req         tx timeout                                0
d-req         tx try limit                              0
d-req         tx failure                                0
d-req         discarded                                 0
d-req         received answer is over limit             0
d-req         tx failure: no memory                     0
d-req         base-app-tx-timeout                       0
d-req         base-app-rx-timeout                       0
d-req         base-app-tx-discard                       0
d-req         base-app-rx-discard                       0
ayt           bad flags                                 0
ayt           bad fixed destination                     0
ayt           bad routed destination                    0
ayt           tx is over limit                          0
ayt           bad end-to-end id                         0
ayt           no peer for tx                            1248
ayt           peer down while waiting for answer        0
ayt           timeout while waiting for answer          0
ayt           tx timeout                                0
ayt           tx try limit                              0
ayt           tx failure                                0
ayt           discarded                                 0
ayt           received answer is over limit             0
ayt           tx failure: no memory                     0
ayt           base-app-tx-timeout                       0
ayt           base-app-rx-timeout                       0
```

```
  ayt           base-app-tx-discard                     0
  ayt           base-app-rx-discard                     2
  awd           bad flags                               0
  awd           bad fixed destination                   0
  awd           bad routed destination                  0
  awd           tx is over limit                        0
  awd           bad end-to-end id                       0
  awd           no peer for tx                          42
  awd           peer down while waiting for answer      0
  awd           timeout while waiting for answer        0
  awd           tx timeout                              0
  awd           tx try limit                            0
  awd           tx failure                              0
  awd           discarded                               0
  awd           received answer is over limit           0
  awd           tx failure: no memory                   0
  awd           base-app-tx-timeout                     0
  awd           base-app-rx-timeout                     0
  awd           base-app-tx-discard                     0
  awd           base-app-rx-discard                     0

Gx-plus upstream counters:
  Category      Counter                                 Value
  prov          start                                   6
  prov          start: failure: bad data                0
  prov          start: ignore                           0
  prov          start: failure: in recovery             0
  prov          start: failure: spruious                0
  prov          start: prov canceled                    0
  prov          start: threshold canceled               0
  prov          start: rule-report canceled             0
  prov          start: rule-report posted to CCR-T      0
  prov          start: no proc required                 0
  prov          start: success + logout dropped         0
  prov          start: failure + logout dropped         0
  prov          start: success: no-config               0
  prov          start: failure: no-config               0
  prov          start: success: over active limit       0
  prov          start: failure: over active limit       0
  prov          start: success: pending                 0
  prov          start: failure: duplicate               0
  prov          start: success: no-memory               0
  prov          start: failure: no-memory               0
```

```
prov          start: success: message posted              0
prov          start: success: message to ready-queue      0
prov          start: async: message posted                4
prov          start: async: message to ready queue        2
prov          message posted from ready queue             2
prov          response: late response                     0
prov          response: positive                          6
prov          response: negative                          0
prov          response: bad data                          0
prov          response timeout                            2
prov          timeout: message already in ready queue     12
prov          timeout: message posted                     0
prov          timeout: message to ready queue             0
prov          active refill: message posted               0
prov          active refill: message to ready queue       0
prov          message build failure                       0
prov          post failure                                0
prov          implied post                                5
prov          discard                                     0
prov          discard(no effect)                          0
prov          implied response                            0
prov          not supported                               0
prov          element allocated                           6
prov          element freed                               6
logout        start                                       6
logout        start: failure: bad data                    0
logout        start: ignore                               0
logout        start: failure: in recovery                 0
logout        start: failure: spruious                    0
logout        start: prov canceled                        0
logout        start: threshold canceled                   0
logout        start: rule-report canceled                 0
logout        start: rule-report posted to CCR-T          0
logout        start: no proc required                     0
logout        start: success + logout dropped             0
logout        start: failure + logout dropped             0
logout        start: success: no-config                   0
logout        start: failure: no-config                   0
logout        start: success: over active limit           0
logout        start: failure: over active limit           0
logout        start: success: pending                     0
logout        start: failure: duplicate                   0
logout        start: success: no-memory                   0
```

```
logout        start: failure: no-memory                  0
logout        start: success: message posted             4
logout        start: success: message to ready-queue     2
logout        start: async: message posted               0
logout        start: async: message to ready queue       0
logout        message posted from ready queue            2
logout        response: late response                    0
logout        response: positive                         6
logout        response: negative                         0
logout        response: bad data                         0
logout        response timeout                           0
logout        timeout: message already in ready queue    12
logout        timeout: message posted                    0
logout        timeout: message to ready queue            0
logout        active refill: message posted              0
logout        active refill: message to ready queue      0
logout        message build failure                      0
logout        post failure                               0
logout        implied post                               0
logout        discard                                    0
logout        discard(no effect)                         0
logout        implied response                           0
logout        not supported                              0
logout        element allocated                          6
logout        element freed                              6
ruleReport    start                                      12
ruleReport    start: failure: bad data                   0
ruleReport    start: ignore                              0
ruleReport    start: failure: in recovery                0
ruleReport    start: failure: spruious                   0
ruleReport    start: prov canceled                       0
ruleReport    start: threshold canceled                  0
ruleReport    start: rule-report canceled                0
ruleReport    start: rule-report posted to CCR-T         6
ruleReport    start: no proc required                    0
ruleReport    start: success + logout dropped            0
ruleReport    start: failure + logout dropped            0
ruleReport    start: success: no-config                  0
ruleReport    start: failure: no-config                  0
ruleReport    start: success: over active limit          0
ruleReport    start: failure: over active limit          0
ruleReport    start: success: pending                    0
ruleReport    start: failure: duplicate                  0
```

```
ruleReport    start: success: no-memory                  0
ruleReport    start: failure: no-memory                  0
ruleReport    start: success: message posted             6
ruleReport    start: success: message to ready-queue     0
ruleReport    start: async: message posted               0
ruleReport    start: async: message to ready queue       0
ruleReport    message posted from ready queue            0
ruleReport    response: late response                    0
ruleReport    response: positive                         6
ruleReport    response: negative                         0
ruleReport    response: bad data                         0
ruleReport    response timeout                           0
ruleReport    timeout: message already in ready queue    0
ruleReport    timeout: message posted                    0
ruleReport    timeout: message to ready queue            0
ruleReport    active refill: message posted              0
ruleReport    active refill: message to ready queue      0
ruleReport    message build failure                      0
ruleReport    post failure                               0
ruleReport    implied post                               0
ruleReport    discard                                    0
ruleReport    discard(no effect)                         0
ruleReport    implied response                           0
ruleReport    not supported                              0
ruleReport    element allocated                          6
ruleReport    element freed                              6
threshold     start                                      113
threshold     start: failure: bad data                   0
threshold     start: ignore                              0
threshold     start: failure: in recovery                0
threshold     start: failure: spruious                   0
threshold     start: prov canceled                       0
threshold     start: threshold canceled                  0
threshold     start: rule-report canceled                0
threshold     start: rule-report posted to CCR-T         0
threshold     start: no proc required                    0
threshold     start: success + logout dropped            0
threshold     start: failure + logout dropped            0
threshold     start: success: no-config                  0
threshold     start: failure: no-config                  0
threshold     start: success: over active limit          0
threshold     start: failure: over active limit          0
threshold     start: success: pending                    0
```

```
threshold     start: failure: duplicate                  0
threshold     start: success: no-memory                  0
threshold     start: failure: no-memory                  0
threshold     start: success: message posted             92
threshold     start: success: message to ready-queue     21
threshold     start: async: message posted               0
threshold     start: async: message to ready queue       0
threshold     message posted from ready queue            0
threshold     response: late response                    0
threshold     response: positive                         92
threshold     response: negative                         0
threshold     response: bad data                         0
threshold     response timeout                           21
threshold     timeout: message already in ready queue    0
threshold     timeout: message posted                    0
threshold     timeout: message to ready queue            0
threshold     active refill: message posted              0
threshold     active refill: message to ready queue      0
threshold     message build failure                      0
threshold     post failure                               0
threshold     implied post                               0
threshold     discard                                    0
threshold     discard(no effect)                         0
threshold     implied response                           0
threshold     not supported                              0
threshold     element allocated                          113
threshold     element freed                              113
general       start                                      0
general       start: failure: bad data                   0
general       start: ignore                              0
general       start: failure: in recovery                0
general       start: failure: spruious                   0
general       start: prov canceled                       0
general       start: threshold canceled                  0
general       start: rule-report canceled                0
general       start: rule-report posted to CCR-T         0
general       start: no proc required                    0
general       start: success + logout dropped            0
general       start: failure + logout dropped            0
general       start: success: no-config                  0
general       start: failure: no-config                  0
general       start: success: over active limit          0
general       start: failure: over active limit          0
```

```
  general       start: success: pending                     0
  general       start: failure: duplicate                   0
  general       start: success: no-memory                   0
  general       start: failure: no-memory                   0
  general       start: success: message posted              0
  general       start: success: message to ready-queue      0
  general       start: async: message posted                0
  general       start: async: message to ready queue        0
  general       message posted from ready queue             0
  general       response: late response                     0
  general       response: positive                          0
  general       response: negative                          0
  general       response: bad data                          0
  general       response timeout                            0
  general       timeout: message already in ready queue     0
  general       timeout: message posted                     0
  general       timeout: message to ready queue             0
  general       active refill: message posted               0
  general       active refill: message to ready queue       0
  general       message build failure                       0
  general       post failure                                0
  general       implied post                                0
  general       discard                                     0
  general       discard(no effect)                          0
  general       implied response                            0
  general       not supported                               0
  general       element allocated                           0
  general       element freed                               0

Gx-plus upstream diameter-events:
  Category      Counter                                     Value
  prov          bad flags                                   0
  prov          bad fixed destination                       0
  prov          bad routed destination                      0
  prov          tx is over limit                            0
  prov          bad end-to-end id                           0
  prov          no peer for tx                              0
  prov          peer down while waiting for answer          0
  prov          timeout while waiting for answer            0
  prov          tx timeout                                  0
  prov          tx try limit                                0
  prov          tx failure                                  0
  prov          discarded                                   0
```

```
prov          received answer is over limit              0
prov          tx failure: no memory                      0
prov          base-app-tx-timeout                        0
prov          base-app-rx-timeout                        0
prov          base-app-tx-discard                        0
prov          base-app-rx-discard                        0
logout        bad flags                                  0
logout        bad fixed destination                      0
logout        bad routed destination                     0
logout        tx is over limit                           0
logout        bad end-to-end id                          0
logout        no peer for tx                             0
logout        peer down while waiting for answer         0
logout        timeout while waiting for answer           0
logout        tx timeout                                 0
logout        tx try limit                               0
logout        tx failure                                 0
logout        discarded                                  0
logout        received answer is over limit              0
logout        tx failure: no memory                      0
logout        base-app-tx-timeout                        0
logout        base-app-rx-timeout                        0
logout        base-app-tx-discard                        0
logout        base-app-rx-discard                        0
ruleReport    bad flags                                  0
ruleReport    bad fixed destination                      0
ruleReport    bad routed destination                     0
ruleReport    tx is over limit                           0
ruleReport    bad end-to-end id                          0
ruleReport    no peer for tx                             0
ruleReport    peer down while waiting for answer         0
ruleReport    timeout while waiting for answer           0
ruleReport    tx timeout                                 0
ruleReport    tx try limit                               0
ruleReport    tx failure                                 0
ruleReport    discarded                                  0
ruleReport    received answer is over limit              0
ruleReport    tx failure: no memory                      0
ruleReport    base-app-tx-timeout                        0
ruleReport    base-app-rx-timeout                        0
ruleReport    base-app-tx-discard                        0
ruleReport    base-app-rx-discard                        0
threshold     bad flags                                  0
```

```
  threshold     bad fixed destination                       0
  threshold     bad routed destination                      0
  threshold     tx is over limit                            0
  threshold     bad end-to-end id                           0
  threshold     no peer for tx                              0
  threshold     peer down while waiting for answer          0
  threshold     timeout while waiting for answer            0
  threshold     tx timeout                                  0
  threshold     tx try limit                                0
  threshold     tx failure                                  0
  threshold     discarded                                   0
  threshold     received answer is over limit               0
  threshold     tx failure: no memory                       0
  threshold     base-app-tx-timeout                         0
  threshold     base-app-rx-timeout                         0
  threshold     base-app-tx-discard                         0
  threshold     base-app-rx-discard                         0
  general       bad flags                                   0
  general       bad fixed destination                       0
  general       bad routed destination                      0
  general       tx is over limit                            0
  general       bad end-to-end id                           0
  general       no peer for tx                              0
  general       peer down while waiting for answer          0
  general       timeout while waiting for answer            0
  general       tx timeout                                  0
  general       tx try limit                                0
  general       tx failure                                  0
  general       discarded                                   0
  general       received answer is over limit               0
  general       tx failure: no memory                       0
  general       base-app-tx-timeout                         0
  general       base-app-rx-timeout                         0
  general       base-app-tx-discard                         0
  general       base-app-rx-discard                         0

Gx-plus downstream counters:
  Category      Counter                                     Value
  audit         new request                                 0
  audit         retry request                               0
  audit         queue full                                  0
  audit         no-memory                                   0
  audit         message posted                              0
```

```
audit          message build failure                    0
audit          post failure                             0
d-exact        new request                              0
d-exact        retry request                            0
d-exact        queue full                               0
d-exact        no-memory                                0
d-exact        message posted                           0
d-exact        message build failure                    0
d-exact        post failure                             0
d-bulk         new request                              5
d-bulk         retry request                            0
d-bulk         queue full                               0
d-bulk         no-memory                                0
d-bulk         message posted                           5
d-bulk         message build failure                    0
d-bulk         post failure                             0
d-done         new request                              0
d-done         retry request                            0
d-done         queue full                               0
d-done         no-memory                                0
d-done         message posted                           0
d-done         message build failure                    0
d-done         post failure                             0
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear network-access gx-plus replay | **2610**

clear network-access gx-plus statistics | **2611**

show network-access nasreq statistics

IN THIS SECTION

- Syntax | 3032
- Description | 3032
- Required Privilege Level | 3032
- Output Fields | 3032
- Sample Output | 3036
- Release Information | 3037

Syntax

```
show network-access nasreq statistics
```

Description

Display information about the NASREQ protocol exchanges that are used to authenticate and authorize subscribers when they log in.

Required Privilege Level

view

Output Fields

Table 152 on page 3033 lists the output fields for the `show network-access nasreq statistics` command. Output fields are listed in the approximate order in which they appear.

Table 152: show network-access nasreq statistics Output Fields

Field Name	Field Description
`Authentication request attempts`	Number of authentication requests forwarded to the NASREQ engine. Note that authentication includes authorization in one message exchange.
`Authentication request failures`	Number of authentication request failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Authentication request messages sent`	Number of authentication requests forwarded to the Diameter engine.
`Authentication request message failures`	Number of authentication requests failures at the Diameter engine.
`Authentication request messages timeouts`	Number of authentication requests that timed out waiting for a response from the NASREQ server.
`Authentication denies failures`	Number of authentication responses with an unsuccessful result-code.
`Authentication grants received`	Number of authentication responses with a successful result-code.
`Authorization request attempts`	Number of authorization requests to the NASREQ engine.
`Authorization request failures`	Number of authorization request failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Authorization request messages sent`	Number of authorization requests forwarded to the Diameter engine.
`Authorization request message failures`	Number of authorization requests failures at the Diameter engine.
`Authorization request messages timeouts`	Number of authorization requests that timed out waiting for a response from the NASREQ server.

Table 152: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
`Authorization denies failures`	Number of authorization responses with an unsuccessful result-code.
`Authorization grants received`	Number of authorization responses with a successful result-code.
`Session-terminate request attempts`	Number of session-terminate requests forwarded to the NASREQ engine.
`Session-terminate request failures`	Number of session-terminate requests failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Session-terminate request messages sent`	Number of session-terminate requests forwarded to the Diameter engine.
`Session-terminate request messages failures`	Number of session-terminate requests failures at the Diameter engine.
`Session-terminate request messages timeouts`	Number of session-terminate requests that timed out waiting for a response from the NASREQ server.
`Session-terminate response message failures`	Number of session-terminate responses with an unsuccessful result-code.
`Session-terminate response messages received`	Number of session-terminate responses with a successful result-code.
`Abort-session requests received`	Number of Abort-Session-Request messages received by the NASREQ engine.
`Abort-session response-ack messages sent`	Number of Abort-Session-Request messages receiving an ACK response.

Table 152: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
Abort-session response-nack messages sent	Number of Abort-Session-Request messages receiving a NACK response.
Abort-session response message failures	Number of Abort-Session-Request messages that failed transmission, either due to timeout in the transmit queue, or failed by the Diameter engine.
Number of NASREQ subscribers	Number of active subscribers that received NASREQ authorization.
Number of result-code-1xxx	Number of response messages with an informational result-code.
Number of result-code-2xxx	Number of response messages with a success result-code.
Number of result-code-3xxx	Number of response messages with a protocol-error result-code.
Number of result-code-4xxx	Number of response messages with an transient-error result-code.
Number of result-code-5xxx	Number of response messages with a permanent-error result-code.
Number of result-code-other	Number of response messages with an unregistered result-code.
Transmit queue time-outs	Number of request/responses that timed out waiting for Diameter engine resources.
Transmit message drops	Number of request/responses dropped.
Transmit message failures	Number of request/response failures at the Diameter engine.
Number of requests on the Tx queue	Number of requests/responses currently on the transmit queue.

Table 152: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
Number of requests waiting for a response	Number of requests currently in the response queue waiting for a response.
Number of outstanding requests	Number of outstanding requests, including AAR and STR.
Total number of allocated messages	Number of currently allocated request/response messages.

Sample Output

command-name

```
user@host> show network-access nasreq statistics
Authentication request attempts:4000
Authentication request failures:0
Authentication request messages sent:4000
Authentication request message failures:0
Authentication request messages timeouts:0
Authentication denies failures:0
Authentication grants received:4000
Authorization request attempts:4000
Authorization request failures:0
Authorization request messages sent:4000
Authorization request message failures:0
Authorization request messages timeouts:0
Authorization denies failures:0
Authorization grants received:4000
Session-terminate request attempts:1000
Session-terminate request failures:0
Session-terminate request messages sent:1000
Session-terminate request message failures:0
Session-terminate request messages timeouts:0
Session-terminate response messages failures:0
Session-terminate response messages received:1000
Abort-session requests received:1000
Abort-session response-ack messages sent:1000
```

```
Abort-session response-nack messages sent:0
Abort-session response message failures:0
Number of NASREQ subscribers:3000
Number of result-code-1xxx:0
Number of result-code-2xxx:5000
Number of result-code-3xxx:0
Number of result-code-4xxx:0
Number of result-code-5xxx:0
Number of result-code-other:0
Transmit queue time-outs:0
Transmit message drops:0
Transmit message failures:0
Number of requests on the Tx queue:0
Number of requests waiting for a response:0
Number of outstanding requests:0
Total number of allocated messages:0
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Diameter Network Access Server Application (NASREQ) | **1463**
- Configuring the Diameter Network Access Server Application (NASREQ) | **1465**
- Messages Used by Diameter Applications | **1342**

show network-access ocs

IN THIS SECTION

- Syntax | **3038**
- Description | **3038**
- Options | **3038**
- Required Privilege Level | **3038**

- Output Fields | 3038
- Sample Output | 3039
- Release Information | 3040

Syntax

```
show network-access ocs
<state>
<statistics [brief | detail | extensive]>
```

Description

Display Online Charging System (OCS) provisioning state and statistics information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
state	(Optional) Display OCS provisioning state.
statistics	(Optional) Display OCS statistics.

Required Privilege Level

view

Output Fields

Table 153 on page 3039 lists the output fields for the `show network-access ocs` command. Output fields are listed in the approximate order in which they appear.

Table 153: show network-access ocs Output Fields

Field Name	Field Description	Level of Output
Ocs state	State of the OCS components, including the following: • **created**—Just created. • **dead**—Ready to be deleted. • **final**—Final interrogation in progress. • **idle**—Idle. • **initial**—Initial interrogation in progress. • **intermediate**—Intermediate interrogation in progress.	All levels
Ocs general counters	Type and number of Diameter response and answer messages for OCS communication using the Gy protocol.	All levels

Sample Output

show network-access ocs state

```
user@host> show network-access ocs state
Ocs state:
  Component                   Value
  state                       active
  active-configuration        yes
  queue-state                 normal
  subscribers                 2000
  sub-in-idle                 2000
```

show network-access ocs statistics

```
user@host> show network-access ocs statistics
Ocs general counters:
  Counter                                                 Value
```

ccr-gy-i	4001
cca-gy-i	4001
ccr-gy-u	38001
cca-gy-u	38001
ccr-gy-t	2001
cca-gy-t	2001
asr	1
asa-ack	1

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- ocs (Diameter Applications) | **2090**
- clear network-access ocs statistics | **2613**
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers | **1413**
- Understanding Gy Interactions Between the Router and the OCS | **1431**
- Understanding Interactions Between the PCRF, PCEF, and OCS | **1439**

show network-access pcrf

IN THIS SECTION

- Syntax | **3041**
- Description | **3041**
- Options | **3041**
- Required Privilege Level | **3041**
- Output Fields | **3041**
- Sample Output | **3044**
- Release Information | **3045**

Syntax

```
show network-access pcrf
<state>
<statistics [brief | detail | extensive]>
<subscribers>
```

Description

Display Policy and Charging Rules Function (PCRF) provisioning state and statistics information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
state	(Optional) Display PCRF provisioning state.
statistics	(Optional) Display PCRF statistics.
subscribers	(Optional) Display session ID and provisioning state for each PCRF subscriber.

Required Privilege Level

view

Output Fields

Table 154 on page 3042 lists the output fields for the `show network-access pcrf` command. Output fields are listed in the approximate order in which they appear.

Table 154: show network-access pcrf Output Fields

Field Name	Field Description	Level of Output
`Pcrf state`	State of the PCRF components: • **active**—Installed services from local grant. • **created**—Just created. • **dead**—Received logout response or timed out waiting for logout response. • **deny**—Received deny from remote and waiting for subscriber logout. • **failed**—GRES before active state or provisioning failed, waiting for subscriber logout. • **grant**—Received remote grant and installing services. • **local-active**—Services installed after local grant. • **local-deny**—Denied by local decision waiting for subscriber logout. • **local-grant**—Received local grant and installing services. • `local-reinit`—Router is waiting for the CCA-GX-T so it can reinitialize the PCRF session. Local reinitialization must be configured on the PCRF partition. This state results from one of the following: ◦ When a RAR is received while the PCRF is in local-active state. ◦ When a RAR is received while the PCRF is in started state and no default service is configured. ◦ When the provisioning-done event occurs while the PCRF is in local-reinit-early state. • `local-reinit-early`—The PCRF is engaged in provisioning services when a failure is detected or the PCRF server sends a RAR message of any type. The session transitions	All levels

Table 154: show network-access pcrf Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	to local-reinit state when the provisioning completes. Local reinitialization must be configured on the PCRF partition. • **logout**—Waiting for logout response. • **ocs-final-wait**—Waiting for OCS to complete final interrogation. • **report**—Sending report. • **update**—Processing RAR-update. • **upd-from-local**—Received remote grant, in local-active state, and installing services. • **started**—Log in.	
`Pcrf general counters`	Type and number of Diameter response and answer messages for PCRF communication using the Gx protocol.	All levels
`Pcrf diameter event counters`	State and number of Diameter events.	All levels
`Pcrf rule install counters`	Result and number of rule installations.	All levels
`Name`	PCRF subscriber name.	All levels
`Session-ID`	Identifier for the PCRF session.	All levels

Sample Output

show network-access pcrf state

```
user@host> show network-access pcrf state
Pcrf state:
  Component                    Value
  state                        active
  active-configuration         yes
  queue-state                  normal
  subscribers                  2000
  sub-in-idle                  2000
```

show network-access pcrf statistics

```
user@host> show network-access pcrf statistics
Pcrf general counters:
  Counter                                                  Value
  ccr-gx-i                                                 4001
  ccr-gx-i-retry                                           4
  cca-gx-i                                                 4001
  ccr-gx-u                                                 4001
  cca-gx-u                                                 4001
  ccr-gx-t                                                 2001
  cca-gx-t                                                 2001

Pcrf diameter event counters:
  Diameter event                                           Value
  no peer for tx                                           4

Pcrf rule install counters:
  Result                                                   Value
  success                                                  16004
  family mismatch                                          5
```

show network-access pcrf subscribers

```
user@host> show network-access pcrf subscribers
Pcrf subscribers:
  Name                          Session-Id          Prov-State
  user45-example                2761                active
  user48-example                12731               local-active
  user51-example                12732               local-active
  user54-example                16887               started
  user57-example                16895               started
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- pcrf (Diameter Applications) | 2174
- clear network-access pcrf | 2614
- Understanding Gx Interactions Between the Router and the PCRF | 1418
- Understanding Interactions Between the PCRF, PCEF, and OCS | 1439
- Understanding Upstream and Downstream Messages for the PCRF | 1444

show network-access requests statistics

IN THIS SECTION

- Syntax | 3046
- Description | 3046
- Required Privilege Level | 3046
- Output Fields | 3046
- show network-access requests statistics | 3046
- Release Information | 3047

Syntax

```
show network-access requests statistics
```

Description

Display authentication statistics for the configured authentication type.

Required Privilege Level

view

Output Fields

Table 155 on page 3046 lists the output fields for the network-access requests statistics command. Output fields are listed in the approximate order in which they appear.

Table 155: show network-access requests statistics Output Fields

Field Name	Field Description
`Total requests received`	Total number of authentication requests that the device received from clients.
`Total responses sent`	Total number of authentication responses that the device sent to the clients.
`Success responses`	Total number of clients that authenticated successfully.
`Failure responses`	Total number of clients that failed to authenticate.

show network-access requests statistics

command-name

```
user@host> show network-access requests statistics
General authentication statistics
  Total requests received: 100
```

```
  Total responses sent: 70
Radius authentication statistics
  Total requests received: 40
  Success responses: 20
  Failure responses: 20
Radius reauthentication statistics
  Total requests received: 0
  Success responses: 0
  Failure responses: 0
LDAP authentication statistics
  Total requests received: 30
  Success responses: 15
  Failure responses: 15
Local authentication statistics
  Total requests received: 5
  Success responses: 2
  Failure responses: 3
Local re-authentication statistics
  Total requests received: 0
  Success responses: 0
  Failure responses: 0
Securid authentication statistics
  Total requests received: 15
  Success responses: 3
  Failure responses: 12
```

Release Information

Command modified in Release 9.1 of Junos OS.

RELATED DOCUMENTATION

clear network-access requests statistics

show network-access s6a

IN THIS SECTION

- Syntax | 3048
- Description | 3048
- Options | 3048
- Required Privilege Level | 3048
- Sample Output | 3049
- Release Information | 3050

Syntax

```
show network-access s6a
<state>
<statistics>
<statistics extensive>
```

Description

Displays information about the s6a protocol exchanges that are used to authenticate subscribers when subscribers log in.

Options

state	Displays the state of information about authentication.
statistics	Displays information about all s6a general counters.
statistics extensive	Displays detail information about all s6a general counters.

Required Privilege Level

view

Sample Output

show network-access s6a state

```
user@host> show network-access s6a state
S6a state:
  Component                                   Value
  active-configuration                        yes
  queue-state                                 normal
  request-count                               0
```

show network-access s6a statistics

```
user@host> show network-access s6a statistics
S6a general counters:
Counter ............Value
aia-grant ..........1
```

show network-access s6a statistics extensive

```
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
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configure S6a Application | **1379**

authentication-order | 1640
partition (s6a) | 2166

show ppp address-pool

IN THIS SECTION

- Syntax | 3051
- Description | 3051
- Options | 3051
- Required Privilege Level | 3052
- Output Fields | 3052
- Sample Output | 3052
- Release Information | 3053

Syntax

```
show ppp address-pool pool-name
<detail>
```

Description

Display PPP address pool information.

Options

pool-name	Address pool name.
detail	(Optional) Display detailed address pool information.

Required Privilege Level

view

Output Fields

Table 156 on page 3052 lists the output fields for the `show ppp address-pool` command. Output fields are listed in the approximate order in which they appear.

Table 156: show ppp address-pool Output Fields

Field Name	Field Description	Level of Output
`Address pool`	Trace address pool code.	All levels
`Address range`	Range of sequentially ordered IP addresses contained in the address pool.	`detail`
`Number of assigned addresses`	Fixed IP address that is to be given to remote users when they dial in. This is a host-only IP address (subnet mask is 255.255.255.255) and is only for single connection receiver profiles.	All levels
`Number of addresses configured`	Number of IP addresses that are available for allocation and used by PPP sessions.	All levels
`Assigned addresses`	Addresses assigned to PPP sessions from the address pool.	`detail`

Sample Output

show ppp address-pool

```
user@host> show ppp address-pool
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
```

```
  Number of assigned addresses: 0
  Number of addresses configured: 10
```

show ppp address-pool detail

```
user@host> show ppp address-pool ppp1 detail
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
  Number of assigned addresses: 2
  Number of addresses configured: 10
  Assigned addresses:
      203.0.113.221
      203.0.113.222
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Verifying and Managing PPP Configuration for Subscriber Management

show static-subscribers sessions

IN THIS SECTION

- Syntax | 3054
- Description | 3054
- Options | 3054
- Required Privilege Level | 3054
- Output Fields | 3054
- Sample Output | 3055
- Release Information | 3057

Syntax

```
show static-subscribers sessions
<group group-name>
<interface interface-name>
```

Description

Display information about the subscriber sessions for all static subscribers, all static subscribers on an interface group, or a single subscriber on an interface.

Options

group-name (Optional) Display session information for static subscribers on all interfaces in the specified group.

interface-name (Optional) Display session information for the static subscriber on the specified in the specified group.

Required Privilege Level

view

Output Fields

Table 157 on page 3054 lists the output fields for the `show static-subscribers sessions` command. Output fields are listed in the approximate order in which they appear.

Table 157: show static-subscribers sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the interface.	None specified

Table 157: show static-subscribers sessions Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the static subscriber session: • authenticating—Subscriber is being authenticated. • activating client—Client is being activated. • activating services—Subscriber services are being activated. • deactivating client—Client is being deactivated. • deactivating services—Subscriber services are being deactivated. • initializing—Process is initializing. • logged in—Subscriber is logged in to the interface. • logged out—Subscriber is logged out of the interface. • processing statistics—Session statistics are being processed. • terminating session—Subscriber session is being terminated.	None specified
Group	Name of the interface group to which the interface belongs.	None specified
User Name	Username used for the static subscriber. Can be the interface name.	None specified

Sample Output

show static-subscribers sessions

```
user@host> show static-subscribers sessions

Static subscriber information:
Interface         State                Group           User Name
ge-9/1/0.1        logged out           SS1             ge-9-1-0.1
ge-9/1/0.10       logged out           SS1             ge-9-1-0.10
ge-9/1/0.100      logged out           SS1             ge-9-1-0.100
ge-9/1/0.11       logged out           SS1             ge-9-1-0.11
```

```
ge-9/1/0.12       logged out             SS1              ge-9-1-0.12
ge-9/1/0.13       logged out             SS1              ge-9-1-0.13
ge-9/1/0.14       logged out             SS1              ge-9-1-0.14
ge-9/1/0.15       logged out             SS1              ge-9-1-0.15
ge-9/1/0.16       logged out             SS1              ge-9-1-0.16
ge-9/1/0.17       logged out             SS1              ge-9-1-0.17
ge-9/1/0.18       logged out             SS1              ge-9-1-0.18
ge-9/1/0.19       logged out             SS1              ge-9-1-0.19
ge-9/1/0.2        logged out             SS1              ge-9-1-0.2
ge-9/1/0.20       logged out             SS1              ge-9-1-0.20
ge-9/1/0.21       logged out             SS1              ge-9-1-0.21
```

show static-subscribers sessions group

```
user@host> show static-subscribers sessions group boston

Interface         State                  Group            User Name
ge-0/0/1.1        logged in              boston           ge-0/0/1.1
ge-0/0/1.2        logged in              boston           ge-0/0/1.2
```

show static-subscribers sessions interface

```
user@host> show static-subscribers sessions interface ge-0/0/1.1

Interface         State                  Group            User Name
ge-0/0/1.1        logged in              foo              ge-0/0/1.1
```

show static-subscribers sessions interface (Pseudowire Interface)

```
user@host> show static-subscribers sessions interface ps100.4040
Static subscriber information:
Interface         State                  Group            User Name
ps100.4040        logged in              foo              static-user:4000@jnpr.net
```

show static-subscribers sessions group (Pseudowire Interface)

```
user@host> show static-subscribers sessions group grp-1
Static subscriber information:
Interface         State                Group           User Name
ps100.4040        logged in            foo             static-user:4000@jnpr.net
ps100.4041        logged in            foo             static-user:4001@jnpr.net
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Subscribers on Static Interfaces Overview | 1483

show subscribers

IN THIS SECTION

- Syntax | 3058
- Description | 3058
- Options | 3058
- Required Privilege Level | 3062
- Output Fields | 3062
- Sample Output | 3074
- Sample Output for AGF | 3108
- Release Information | 3109

Syntax

```
show subscribers
<detail | extensive | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier>
<agent-remote-identifier agent-remote-identifier>
<aggregation-interface-set-name interface-set-name>
<client-type client-type>
<count>
<id session-id <accounting-statistics>>
<interface interface <accounting-statistics>>
<logical-system logical-system>
<mac-address mac-address>
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<sub-system (agf | bng)>
<subscriber-state subscriber-state>
<user-name user-name>
<vci vci-identifier>
<vpi vpi-identifier>
<vlan-id vlan-id>
```

Description

Display information for active subscribers on both the broadband network gateway (BNG) and Access Gateway Function (AGF) subsystems. When you have subscribers logged in to both the BNG and the AGF, you can display the outputs specific to either subsytem with the `sub-system` option.

See "Sample Output for AGF" on page 3108 for subscriber session information about the AGF subsystem.

Options

detail | extensive | terse	(Optional) Display the specified level of output.

aci-interface-set-name

(Optional) Display all the dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. You must use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) control packets.

address

(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid and rejects the command.

agent-circuit-identifier

(Optional) Display all the dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

Junos OS Release	Substring Support
Junos OS Release 13.3R1	You can specify a substring without a wildcard.
Starting in Junos OS Release 14.1R1	You must specify the complete ACI string; you cannot specify a wildcard.
Starting in Junos OS Releases 15.1R7, 16.1R7, 16.2R3, 17.1R3, 17.2R3, 17.3R3, 17.4R2, 18.1R2, 18.2R1	You can specify a substring, but you must include the wildcard character at the end of the substring.

agent-remote-identifier

(Optional) Display all the dynamic subscriber sessions whose agent remote identifier (ARI) value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.

aggregation-interface-set-name *interface-set-name*

(Optional) Display summary information for the specified aggregation node interface set, including interface, VLAN ID, username, and logical system and routing instance (LS:RI).

client-type (Optional) Display subscribers whose client type matches one of the following client types:

- `dhcp`—Dynamic Host Configuration Protocol (DHCP) clients only.
- `dotlx`—802.1X clients only.
- `essm`—Extensible Subscribers Services Manager (ESSM) clients only.
- `fixed-wireless-access`—Fixed wireless access clients only.
- `fwauth`—FwAuth (authenticated across a firewall) clients only.
- `l2tp`—Layer 2 Tunneling Protocol (L2TP) clients only.
- `mlppp`—Multilink Point-to-Point Protocol (MLPPP) clients only.
- `ppp`—Point-to-Point Protocol (PPP) clients only.
- `pppoe`—Point-to-Point Protocol over Ethernet (PPPoE) clients only.
- `static`—Static clients only.
- `vlan`—VLAN clients only.
- `vlan-oob`—VLAN out-of-band (ANCP-triggered) clients only.
- `vpls-pw`—Virtual private LAN service (VPLS) pseudowire clients only.
- `xauth`—Xauth clients only.

count (Optional) Display the count of the total subscribers and active subscribers for any specified option. You can use the `count` option alone or with the `address`, `client-type`, `interface`, `logical-system`, `mac-address`, `profile-name`, `routing-instance`, `stacked-vlan-id`, `subscriber-state`, or `vlan-id` options.

id *session-id* (Optional) Display a specific subscriber session whose session ID matches the specified subscriber ID. You can display the subscriber IDs by using the `show subscribers extensive` command or the `show subscribers interface extensive` command.

id *session-id* accounting-statistics (Optional) Display accurate subscriber accounting statistics for a subscriber session with the specified ID. Requires the `actual-transmit-statistics` statement to be configured in the dynamic profile for the dynamic logical interface. If the statement is not configured, a value of 0 is displayed for accounting statistics.

interface (Optional) Display subscribers whose interface matches the specified interface.

interface **accounting-statistics**	(Optional) Display subscriber accounting statistics for the specified interface. Requires the `actual-transmit-statistics` statement to be configured in the dynamic profile for the dynamic logical interface.
logical-system	(Optional) Display subscribers whose logical system matches the specified logical system.
mac-address	(Optional) Display subscribers whose MAC address matches the specified MAC address.
physical-interface-name	(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.
profile-name	(Optional) Display subscribers whose dynamic profile matches the specified profile name.
routing-instance	(Optional) Display subscribers whose routing instance matches the specified routing instance.
stacked-vlan-id	(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.
subsystem (agf | bng)	(Optional) Display subscriber information on the AGF or BNG subsystem. **NOTE**: The subsystem option is only available when both the AGF and BNG subscribers are logged into the router at the same time.
subscriber-state	(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).
user-name	(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.
vci-identifier	(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active Asynchronous Transfer Mode (ATM) subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.
vpi-identifier	(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id (Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the `stacked-vlan-id` *stacked-vlan-id* option to match the outer VLAN tag.

NOTE: Because of display limitation, the logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 158 on page 3062 lists the output fields for the `show subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 158: show subscribers Output Fields

Field Name	Field Description
`Interface`	Interface associated with the subscriber. The router or the switch displays the subscribers whose interface matches or begins with the specified interface. The asterisk (*) indicates a continuation of addresses for the same session.
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions, the value is `Tunnel-switched`.
`User Name`	Name of subscriber.
`LS:RI`	Logical system and routing instance associated with the subscriber.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Type`	Subscriber client type (DHCP, FWA, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
`IP Address`	Subscriber IPv4 address.
`IP Netmask`	Subscriber IP netmask. (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the L2TP network server (LNS) generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
`Primary DNS Address`	IP address of the primary Domain Name System (DNS) server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`Secondary DNS Address`	IP address of the secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Primary DNS Address`	IPv6 address of the primary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Secondary DNS Address`	IPv6 address of the secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Domain name server inet`	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Domain name server inet6`	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Primary WINS Address`	IP address of the primary Windows Internet Name Service (WINS) server.
`Secondary WINS Address`	IP address of the secondary WINS server.
`IPv6 Address`	Subscriber IPv6 address, or multiple addresses.
`IPv6 Prefix`	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
`IPv6 User Prefix`	IPv6 prefix obtained through Neighbor Discovery Router Advertisement (NDRA).
`IPv6 Address Pool`	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
`IPv6 Network Prefix Length`	Length of the network portion of the IPv6 address.
`IPv6 Prefix Length`	Length of the subscriber IPv6 prefix.
`Logical System`	Logical system associated with the subscriber.
`Routing Instance`	Routing instance associated with the subscriber.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Interface`	(Enhanced subscriber management for the MX Series routers) Name of the enhanced subscriber management logical interface, in the form `demux0.`*nnnn* (for example, `demux0.3221225472`), to which access-internal and framed subscriber routes are mapped.
`Interface Type`	Whether the subscriber interface is `Static` or `Dynamic`.
`Interface Set`	Internally generated name of the dynamic agent circuit identifier (ACI) or ATM line interface (ALI) interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line. The name of the interface set uses one of the following prefixes: • `aci`—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • `ari`—ARI or agent remote identifier; for example, ari-1033-demux0.3221225524. • `aci+ari`—Both the ACI and the ARI; for example, aci+ari-1033-demux0.3221225524. • `noids`—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. **NOTE**: ACI interface sets are configured with the `agent-circuit-identifier` autoconfiguration stanza. ALI interface sets are configured with the `line-identity` autoconfiguration stanza. Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable $junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.
`Interface Set Type`	Interface type of the ACI interface set: `Dynamic`. This is the only ACI interface set type currently supported.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Interface Set Session ID`	Identifier of the dynamic ACI interface set entry in the session database.
`Underlying Interface`	Name of the underlying interface for the subscriber session.
`Dynamic Profile Name`	Dynamic profile used for the subscriber.
`Dynamic Profile Version`	Version number of the dynamic profile used for the subscriber.
`MAC Address`	MAC address associated with the subscriber.
`State`	Current state of the subscriber session (`Init`, `Configured`, `Active`, `Terminating`, or `Tunneled`).
`L2TP State`	Current state of the L2TP session (`Tunneled` or `Tunnel-switched`). When the value is `Tunnel-switched`, two entries are displayed for the subscriber; the first entry is at the L2TP network server (LNS) interface on the L2TP tunnel switch (LTS) and the second entry is at the L2TP access concentrator (LAC) interface on the LTS.
`Tunnel switch Profile Name`	Name of the L2TP tunnel switch profile that initiates tunnel switching.
`Local IP Address`	IP address of the local gateway (LAC).
`Remote IP Address`	IP address of the remote peer (LNS).
`PFE Flow ID`	Forwarding flow identifier.
`VLAN Id`	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
`Stacked VLAN Id`	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	For the dhcp client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in the hexadecimal format. For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSL access multiplexer (DSLAM) interface on which the subscriber request originates.
Agent Remote ID	For the dhcp client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the network access server (NAS)-facing DSLAM interface on which the subscriber request originates.
Aggregation Interface-set Name	Value of the $junos-aggregation-interface-set-name predefined variable; one of the following: • When the hierarchical-access-network-detection option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the Access Node Control Protocol (ANCP) Port Up message or PPPoE Active Discover Request (PADR) IA tags begins with a # character, then the variable takes the value of the remainder of the string after the # character. • When the hierarchical-access-network-detection option is not configured, or if the sting does not begin with the # character, then the variable takes the value specified with the predefined-variable-defaults statement.
Accounting Statistics	Actual transmitted subscriber accounting statistics by the session ID or the interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the actual-transmit-statistics statement is included in the dynamic profile.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber in kilobits per second.
IPv4 Input Service Set	Input service set in access dynamic profile.
IPv4 Output Service Set	Output service set in access dynamic profile.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
PCEF Profile	Policy and charging enforcement function (PCEF) profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to the IPv4 packets that fail the reverse-path-forwarding (RPF) check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to the IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS Change of Authorization or CoA) associated with the subscribers.
Service Session ID	ID number for a subscriber service session.
Service Session Name	Service session profile name.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds for which the subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable $junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).
DSL type	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: ADSL, ADSL2, ADSL2+, OTHER, SDSL, VDSL, or VDSL2.
Frame/Cell Mode	Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091): • Cell—When the DSL line type is one of the following: asymmetric digital subscriber line (ADSL), asymmetric digital subscriber line 2 (ADSL2), or asymmetric digital subscriber line 2 plus (ADSL2+). • Frame—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. The value is stored in the subscriber session database.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Overhead accounting bytes`	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
`Actual upstream data rate`	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).
`Actual downstream data rate`	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
`Adjusted downstream data rate`	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
`Adjusted upstream data rate`	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database. **NOTE**: This output field is only available on Junos Release 19.1R1 or earlier versions.
`AGF Mode`	Type of access. Adaptive mode indicates an FN-RG.
`Local TEID-U`	Tunnel endpoint identifier on the BNG for the GTP-U user plane tunnel to the eNodeB. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPU Tunnel Local IP address` value.
`Local TEID-C`	Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the MME. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPC Local IP address` value.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Remote TEID-U`	Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB. A fully qualified remote TEID-U consists of this identifier and the `GTPU Tunnel Remote IP address` value.
`Remote TEID-C`	Tunnel endpoint identifier on the MME for the GPRS tunneling protocol, control (GTP-C) plane tunnel to the BNG. The identifier is allocated by the MME. A fully qualified remote TEID-C consists of this identifier and the `GTPC Remote IP address` value.
`GTPU Tunnel Remote IP address`	IP address of the S1-U interface on the eNodeB for the GPRS tunneling protocol, user plane (GTP-U) tunnel endpoint. A fully qualified remote TEID-U consists of this address and the `Remote TEID-U` value.
`GTPU Tunnel Local IP address`	IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint. A fully qualified local TEID-U consists of this address and the `Local TEID-U` value
`GTPC Remote IP address`	IP address of the S11 interface on the MME for the GTP-C tunnel endpoint. A fully qualified remote TEID-C consists of this address and the `Remote TEID-C` value.
`GTPC Local IP address`	IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint. A fully qualified local TEID-C consists of this address and the `Local TEID-C` value.
`Access Point Name`	Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.
`Tenant`	Name of the tenant system. You can create multiple tenant system administrators for a tenant system with different permission levels based on your requirements.

Table 158: show subscribers Output Fields *(Continued)*

Field Name	Field Description
Routing instance	Name of the routing instance. When a custom routing instance is created for a tenant system, all the interfaces defined in that tenant system are added to that routing instance.
Dynamic Profile Version Alias	Configured name for a specific variation of a base dynamic profile. IT's presence indicates that the profile configuration is different from that of the base profile. The value is conveyed to the RADIUS server during authentication in the Client-Profile-Name VSA (26–4874–174).

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
Interface               IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824     10                                     default:default
demux0.1073741824       203.0.113.10        WHOLESALER-CLIENT default:default
demux0.1073741825       203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826       203.0.113.3         RETAILER2-CLIENT  test1:retailer2
```

show subscribers (IPv6)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name           LS:RI
ge-1/0/0.0         2001:db8:c0:0:0:0/74    WHOLESALER-CLIENT   default:default
     *             2001:db8:1/128          subscriber-25       default:default
```

show subscribers (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers
Interface           IP Address/VLAN ID         User Name                      LS:RI
demux0.1073741834   0x8100.1002 0x8100.1                                  default:default
demux0.1073741835   0x8100.1001 0x8100.1                                  default:default
pp0.1073741836      203.0.113.13            dualstackuser1@example1.com     default:ASP-1
```

```
*                   2001:db8:1::/48
*                   2001:db8:1:1::/64
pp0.1073741837      203.0.113.33            dualstackuser2@example1.com     default:ASP-1
*                   2001:db8:1:2:5::/64
```

show subscribers (Single Session DHCP Dual Stack)

```
user@host> show subscribers

Interface           IP Address/VLAN ID        User Name               LS:RI
demux0.1073741364   192.168.10.10             dual-stack-retail35     default:default
                    2001:db8::100:0:0:0/74                            default:default
                    2001:db8:3ffe:0:4::/64
```

show subscribers (Single Session DHCP Dual Stack detail)

```
user@host> show subscribers id 27 detail
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
```

```
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00
```

show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name         LS:RI
si-4/0/0.1         192.0.2.0          user@example.com   default:default
```

show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface           IP Address/VLAN ID    User Name              LS:RI
si-2/1/0.1073741842 Tunnel-switched       user@example.com        default:default
si-2/1/0.1073741843 Tunnel-switched       user@example.com        default:default
```

show subscribers aggregation-interface-set-name

```
user@host> show subscribers aggregation-interface-set-name FRA*

Interface              IP Address/VLAN ID              User Name                 LS:RI
ge-1/0/0.3221225472    50                                ancp                      default:isp1-
subscriber
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
```

```
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 203.0.113.27
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :2560
Login Time: 2009-08-25 14:43:56 PDT
```

show subscribers client-type dhcp detail (DHCPv6)

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
User Name: DEFAULTUSER
IPv6 Address: 2001:db8::2
IPv6 Prefix: 2001:db8:1::/64
Logical System: default
Routing Instance: default
Interface: demux0.3221225602
Interface type: Static
Underlying Interface: demux0.3221225602
Dynamic Profile Name: client-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 142
Session ID: 142
PFE Flow ID: 148
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2018-03-29 12:27:38 EDT
DHCP Options: len 56
00 08 00 02 00 00 00 01 00 0e 00 01 00 01 22 4f d0 33 00 11
```

```
01 00 00 01 00 03 00 0c 00 00 00 0a 00 04 9d 40 00 07 62 00
00 19 00 0c 00 00 00 0b 00 04 9d 40 00 07 62 00
Server DHCPV6 Options: len 94
00 0a 00 06 11 22 33 44 55 66 00 11 00 09 00 00 0c 4c 00 02
00 01 aa 00 11 00 20 00 00 0a 4c 00 02 00 02 32 33 00 03 00
03 34 35 36 00 05 00 06 31 32 33 34 35 36 00 06 00 01 31 00
11 00 09 00 00 0b 4c 00 02 00 01 bb 00 11 00 12 00 00 0d e9
00 01 00 03 aa bb cc 00 02 00 03 dd ee cc
DHCPV6 Header: len 4
01 fc e4 96
```

show subscribers client-type dhcp extensive

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
User Name: user
IP Address: 192.0.2.4
IP Netmask: 255.0.0.0
IPv6 Address: 2001:db8:3::103
IPv6 Prefix: 2001:db8::/68
Domain name server inet6: 2001:db8:1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:00:5e:00:53:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
```

```
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8
00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool
```

show subscribers client-type fixed-wireless-access

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

Interface             IP Address/VLAN ID            User Name                LS:RI
ps1.3221225472        192.0.2.10                    505024101215074          default:default
ps1.3221225473        192.0.2.11                    505024101215075          default:default
```

show subscribers client-type fixed-wireless-access detail (Detail)

```
user@host> show subscribers client-type fixed-wireless-access detail
Type: FWA
User Name: 505024101215074
IP Address: 192.0.2.10
IP Netmask: 255.255.0.0
Interface: ps1.3221225472
Interface type: Dynamic
Dynamic Profile Name: fwa-profile
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 11
Login Time: 2019-04-10 14:10:12 PDT
Local TEID-U: 1
Local TEID-C: 1
Remote TEID-U: 2000000
Remote TEID-C: 1000000
GTPU Tunnel Remote IP Address: 203.0.113.1.3
GTPU Tunnel Local IP Address: 203.0.113.2.5
GTPC Remote IP Address: 203.0.113.1.2
```

```
GTPC Local IP Address: 203.0.113.1.1
Access Point Name: user21
```

show subscribers client-type vlan-oob detail

```
user@host> show subscribers client-type vlan-oob detail
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISP1
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 1234
Session ID: 77
VLAN Id: 126
Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT
```

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 203.0.113.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
IP Netmask: 255.0.0.0
```

```
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example demux0.8201:6544
Session ID: 6544
Agent Circuit ID: ifl3720
Agent Remote ID: ifl3720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: example :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
```

```
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (pseudowire Interface for GRE Tunnel)

```
user@host> show subscribers detail
Interface           IP Address/VLAN ID        User Name           LS:RI
ps0.3221225484      192.0.2.2
ps0.3221225485      192.0.2.3
demux0.3221225486    1                                          default:default
demux0.3221225487    1                                          default:default
demux0.3221225488   198.51.0.1                                  default:default
demux0.3221225489   198.51.0.2                                  default:default
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
```

```
User Name: user@example.com
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
```

```
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.31
Remote IP Address: 192.0.2.1
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
```

```
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.15
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers detail (Dynamic Profile Version Alias)

```
user@host> show subscribers detail

Type: PPPoE
User Name: DEFAULTUSER
IP Address: 192.0.2.21
IP Netmask: 255.255.255.255
IPv6 Address: 2001:db8::17
Logical System: default
Routing Instance: default
Interface: pp0.3221225720
Interface type: Dynamic
Underlying Interface: demux0.3221225719
Dynamic Profile Name: pppoe-client-profile
Dynamic Profile Version Alias: profile-version1a
MAC Address: 00:00:5E:00:53:38
State: Active
Radius Accounting ID: 288
Session ID: 288
PFE Flow ID: 344
VLAN Id: 1
Login Time: 2019-09-23 10:40:56 IST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: uer@host
IP Address: 192.0.2.136
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 15
Session ID: 15
```

```
PFE Flow ID: 2
VLAN Id: 100
Login Time: 2021-05-24 11:30:07 IST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 31 2f
31 2d 30 2d 30 37 05 01 06 0f 21 2c
DHCP Header: len 44
01 01 06 00 00 00 00 1d 00 00 80 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 10 94 00 00 01 00 00 00 00 00 00
00 00 00 00
IP Address Pool: al_pool30
Access Line Attributes:
  Actual upstream data rate: 19998
  Actual downstream data rate: 79999
  Access loop encapsulation: 01 02 00
```

show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes)

```
user@host> show subscribers extensive
Type: VLAN-OOB
User Name: ancp
Logical System: default
Routing Instance: isp1-subscriber
Interface: ge-1/0/0.3221225472
Interface type: Dynamic
Interface Set: FRA-DPU-C-100
Underlying Interface: ge-1/0/0
Core IFL Name: ge-1/0/4.0
Dynamic Profile Name: Prof_L2BSA
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 13
VLAN Id: 50
VLAN Map Id: 20
Inner VLAN Map Id: 1
Inner VLAN Tag Protocol Id: 0x88a8
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
```

```
Login Time: 2018-05-29 08:43:42 EDT
Accounting interval: 72000
Dynamic configuration:
  junos-cos-scheduler-map: 100m
  junos-inner-vlan-tag-protocol-id: 0x88a8
  junos-vlan-map-id: 20

Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225474
Interface type: Dynamic
Interface Set: ge-1/0/0
Underlying Interface: demux0.3221225473
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:10:94:00:00:03
State: Active
Radius Accounting ID: 3
Session ID: 3
PFE Flow ID: 16
Stacked VLAN Id: 50
VLAN Id: 7
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:45 EDT
IP Address Pool: pool-1
Accounting interval: 72000
DSL type: G.fast
Frame/cell mode: Frame
Overhead accounting bytes: 10
Actual upstream data rate: 100000 kbps
Actual downstream data rate: 200000 kbps
Calculated downstream data rate: 180000 kbps
Calculated upstream data rate: 90000 kbps

Adjusted downstream data rate: 160000 kbps
DSL Line Attributes
  Agent Circuit ID: circuit 201
  Agent Remote ID: remote-id
  Actual upstream data rate: 100000
```

```
  Actual downstream data rate: 200000
  DSL type: G.fast
  Access Aggregation Circuit ID: #FRA-DPU-C-100
  Attribute type: 0xAA, Attribute length: 4
    198 51 100 78
```

show subscribers extensive (Passive Optical Network Circuit Interface Set)

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic
Interface Set: otl01.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: user :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|otl01.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: POOL-V4
```

show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
```

```
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (DNS Addresses from RADIUS)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
```

```
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user@example.com
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2016:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
```

```
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48
```

show subscribers extensive (ADF Rules)

```
user@host> show subscribers extensive
...
   Service Session ID: 12
   Service Session Name: SERVICE-PROFILE
   State: Active
   Family: inet
     ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
                       Rule 0: 010101000b0101020b020200201811
                               from {
                                   source-address 203.0.113.232;
                                   destination-address 198.51.100.0/24;
                                   protocol 17;
                               }
                               then {
                                   accept;
                               }
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
```

```
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...
```

show subscribers extensive (PPPoE Subscriber Access Line Rates)

```
user@host> show subscribers extensive
Type: PPPoE
 IP Address: 198.51.100.1
 IP Netmask: 255.255.255.255
 Logical System: default
 Routing Instance: default
 Interface: pp0.3221225475
 Interface type: Dynamic
 Underlying Interface: demux0.3221225474
 Dynamic Profile Name: pppoe-client-profile-with-cos
 MAC Address: 00:00:5e:00:53:02
 State: Active
 Radius Accounting ID: 4
 Session ID: 4
 PFE Flow ID: 14
 Stacked VLAN Id: 40
 VLAN Id: 1
 Agent Circuit ID: circuit0
 Agent Remote ID: remote0
 Login Time: 2017-04-06 15:52:32 PDT

 User Name: DAVE-L2BSA-SERVICE
 Logical System: default
 Routing Instance: isp-1-subscriber
 Interface: ge-1/2/4.3221225472
 Interface type: Dynamic
```

```
  Interface Set: ge-1/2/4
  Underlying Interface: ge-1/2/4
  Core IFL Name: ge-1/3/4.0
  Dynamic Profile Name: L2BSA-88a8-400LL1300VO
  State: Active
Radius Accounting ID: 1
  Session ID: 1
  PFE Flow ID: 14
  VLAN Id: 13
  VLAN Map Id: 102
  Inner VLAN Map Id: 1
  Agent Circuit ID: circuit-aci-3
  Agent Remote ID: remote49-3
  Login Time: 2017-04-05 16:59:29 EDT
  Service Sessions: 4
  IFL Input Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-in
  IFL Output Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-out
  Accounting interval: 900
  DSL type: VDSL
  Frame/Cell Mode: Frame
  Overhead accounting bytes: -10
  Actual upstream data rate: 1024 kbps
  Actual downstream data rate: 4096 kbps
  Adjusted downstream data rate: 3686 kbps
   Dynamic configuration:
    junos-vlan-map-id: 102
    Service Session ID: 5
    Service Session Name: SRL-L1
    State: Active
    Family: inet, inet6
    IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in
    IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out
    Service Activation time: 2017-04-05 16:59:30 EDT
    Dynamic configuration:
      l2bsa-fwf-in: L2BSA-FWF-in-10048
      l2bsa-fwf-out: L2BSA-FWF-out-25088
      rldown: 25088
      rlup: 10048
```

show subscribers extensive (Subscriber Session Using PCEF Profile)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 192.0.2.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
```

```
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default

   Service Session ID: 61
   Service Session Name: pcef-serv-prof
   State: Active
   Family: inet
   IPv4 Input Service Set: tdf-service-set
   IPv4 Output Service Set: tdf-service-set
   PCEF Profile: pcef-prof-1
   PCEF Rule/Rulebase: limit-fb
   Service Activation time: 2017-03-28 08:31:19 PDT
   Dynamic configuration:
     pcef-prof: pcef-prof-1
     pcef-rule1: limit-fb
     svc-filt: svc-filt-1
     svc-set: tdf-service-set
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
```

```
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address:
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
```

```
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:52
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers id accounting-statistics

```
user@host> show subscribers id 601 accounting-statistics
Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

show subscribers interface accounting-statistics

```
user@host> show subscribers interface pp0.3221226949 accounting-statistics
Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
```

```
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322
Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 12
```

```
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: user@test.example.com
IP Address: 192.0.2.0
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface                       IP Address/VLAN ID   User Name         LS:RI
demux0.1073741825           203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826           203.0.113.4         RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 203.0.113.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface           IP Address                              User Name
```

```
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
```

```
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers address detail (Enhanced Subscriber Management)

```
user@host> show subscribers address 203.0.113.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show subscribers extensive (Tenant Systems)

```
user@host:TSYS1> show subscribers extensive
Type: XAUTH
User Name: userX
+    Tenant: TSYS1
    Routing Instance: TSYS1-ri
IP Address: 192.0.2.0
IP Netmask: 203.0.113.0
Primary DNS Address: 198.51.100.0
Secondary DNS Address: 198.51.100.1
Dynamic Profile Name: radius
```

```
State: Active
Session ID: 1
Login Time: 2018-09-18 13:49:00 PDT
```

Sample Output for AGF

The following sample output shows subscribers connected to the AGF:

show subscribers extensive (DHCP on an FN-RG)

```
user@host> show subscribers extensive
Type: DHCP
User Name: USER2
IP Address: 172.16.0.227
Logical System: default
Routing Instance: default
Interface: demux0.3221230587
Interface type: Dynamic
Underlying Interface: demux0.3221230586
Dynamic Profile Name: dhcp-profile
MAC Address: 00:44:46:44:44:44
State: Active
DHCP Relay IP Address: 10.1.0.1
Radius Accounting ID: 5128
Session ID: 5128
PFE Flow ID: 5185
Stacked VLAN Id: 1
VLAN Id: 1
Agent Circuit ID: aci1
Agent Remote ID: ari1
Login Time: 2022-04-26 09:24:56 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
DHCP Header: len 44
01 01 06 00 10 bd b4 93 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 44 46 44 44 44 00 00 00 00 00 00
00 00 00 00
AGF Mode: Adaptive
Local TEID-U: 10183810
Remote TEID-U: 6354992
```

```
GTPU Tunnel Remote IP Address: 10.0.0.1
GTPU Tunnel Local IP Address: 10.0.0.1
5G-QFI: 1
IPv6 Framed Interface Id: 0:7fff:ffff:eaa0
IPv4 Input Filter Name: RG-LWAC-V4-INPUT-FILTER-01-demux0.3221230587-in
IPv4 Output Filter Name: RG-LWAC-V4-OUTPUT-FILTER-01-demux0.3221230587-out
Access Line Attributes:
  Agent Circuit ID: aci1
  Agent Remote ID: ari1
Dynamic configuration:
  junos-cos-guaranteed-rate: 1000000
  junos-cos-guaranteed-rate-burst: 250000
      junos-cos-scheduler: GOLD
          junos-cos-scheduler-shaping-rate: 200000000
          junos-cos-scheduler-tx: 200000000
  junos-cos-scheduler-map: DATA_VOICE_VIDEO_SMAP_UID1468
  junos-cos-shaping-rate: 200000000
  junos-cos-shaping-rate-burst: 50000
  junos-cos-traffic-control-profile: TRAFFIC_CONTROL_PROFILE
  junos-input-filter: RG-LWAC-V4-INPUT-FILTER-01
  junos-input-ipv6-filter: RG-LWAC-V6-INPUT-FILTER-01
  junos-output-filter: RG-LWAC-V4-OUTPUT-FILTER-01
  junos-output-ipv6-filter: RG-LWAC-V6-OUTPUT-FILTER-01

   Service Session ID: 5129
   Service Session Name: SERVICE-PROFILE-BASIC-POLICER
   State: Active
   Family: inet
   Service session type: Service-Profile
   IPv4 Input Filter Name: CAP-POLICER-demux0.3221230587-in
   IPv4 Output Filter Name: CAP-POLICER-demux0.3221230587-out
   Service Activation time: 2022-04-26 09:24:57 PDT
   Dynamic configuration:
     bandwidth-limit: 200k
     burst-size-limit: 75k
```

Release Information

Command introduced in Junos OS Release 9.3.

`client-type`, `mac-address`, `subscriber-state`, and `extensive` options introduced in Junos OS Release 10.2.

`count` option usage with other options introduced in Junos OS Release 10.2.

Options `aci-interface-set-name` and `agent-circuit-identifier` introduced in Junos OS Release 12.2.

The `physical-interface` and `user-name` options introduced in Junos OS Release 12.3.

Options `vci` and `vpi` introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

Options `vci` and `vpi` supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Enhanced subscriber management supported in Junos OS Release 15.1R3 on MX Series routers.

`accounting-statistics` option added in Junos OS Release 15.1R3 and 17.4R1 on MX Series routers.

`aggregation-interface-set-name` option added in Junos OS Release 18.4R1 on MX Series routers.

`sub-system` added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

Verifying and Managing Junos OS Enhanced Subscriber Management

show subscribers summary

IN THIS SECTION

- Syntax | **3111**
- Description | **3111**
- Options | **3111**
- Required Privilege Level | **3112**
- Output Fields | **3112**
- Sample Output | **3115**
- Release Information | **3120**

Syntax

```
show subscribers summary
<all>
<detail | extensive | terse>
<count>
<physical-interface physical-interface-name>
<logical-system logical-system pic | port | routing-instance routing-instance | slot>
<sub-system (agf | bng)>
```

Description

Display summary information for subscribers.

Options

none	Display summary information by state and client type for all subscribers.
all	(Optional) Display summary information by state, client type, and logical system and routing instance (LS:RI).
detail | extensive | terse	(Not supported on MX Series routers) (Optional) Display the specified level of output.
count	(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.
logical-system *logical-system*	(Optional) Display subscribers whose logical system matches the specified logical system.
physical-interface *physical-interface-name*	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type, and LS:RI.
pic	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by PIC number and the total number of subscribers.
port	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by port number and the total number of subscribers.

routing-instance *routing-instance*	(Optional) Display subscribers whose routing instance matches the specified routing instance.
slot	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by Flexible PIC Concentrator (FPC) slot number and the total number of subscribers.
sub-system (agf | bng)	(Optional) Display a count of subscribers using either the Access Gateway Function (AGF) or broadband network gateway (BNG) services.

NOTE: The subsystem option is only available when both the AGF and BNG subscribers are logged into the router at the same time.

NOTE: Because of display limitation, the logical system and routing instance output values are truncated when necessary.

Starting from Junos OS 20.4R1 release, you need a license to use the Extensible Subscribers Services Manager (ESSM) feature.

Required Privilege Level

view

Output Fields

Table 159 on page 3113 lists the output fields for the `show subscribers summary` command. Output fields are listed in the approximate order in which they appear.

Table 159: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
`Subscribers by State`	Number of subscribers summarized by state. The summary includes the following information: • Init—Number of subscribers currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.	`detail` none
`Subscribers by Client Type`	Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN, and VLAN-OOB. This field also displays the total number of subscribers for all client types (Total).	`detailextensive` none
`Subscribers by LS:RI`	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. This field also displays the total number of subscribers for all the LS:RI combinations (Total).	`detail` none
`Subscribers by Connection Type`	Number of subscribers summarized by connection type, `Cross-connected` or `Terminated`.	`extensive`
`Subscribers by Sub-System`	Number of subscribers on the subsystem—AGF or BNG as well as the total number of subscriber.	All levels

Table 159: show subscribers summary Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Interface`	Interface associated with the subscriber. The router or the switch displays subscribers whose interface matches or begins with the specified interface.
	
	The asterisk (*) indicates a continuation of addresses for the same session.
	
	For aggregated Ethernet interfaces, the output of the `summary (pic | port | slot)` options prefixes the interface name with ae0:.
	
	For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and the redundant logical tunnel (RLT) anchor interface. For example:
	
	`ps0: lt-2/1/0`
	`ps1: rlt0: lt-4/0/0`
		All levels

`Count`	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the `summary` option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link. Multiple pseudowire interfaces can share a given logical tunnel or redundant logical tunnel anchor interface. Starting in Junos OS Release 18.1R1, the field displays subscribers per individual pseudowire interface. In earlier releases, the field displays the same number of subscribers for all the pseudowire interfaces that share the same tunnel interface as their anchor point.	`detailextensive` none
`Total Subscribers`	Total number of subscribers for all physical interfaces, all PICs, all ports, or all LS:RI slots.	`detailextensive` none
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*	`terse`

Table 159: show subscribers summary Output Fields *(Continued)*

Field Name	Field Description	Level of Output
User Name	Name of the subscriber.	terse
LS:RI	Logical system and routing instance associated with the subscriber.	terse

Sample Output

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
   Active: 52194
   Total: 52194

Subscribers by Client Type
   DHCP: 10000
   VLAN: 15997
   VLAN-OOB: 3600
   PPPoE: 15998
   ESSM: 6599
   Total: 52194
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
    Init                    3
    Configured             2
    Active                  183
    Terminating             2
    Terminated             1

    TOTAL                   191
```

```
Subscribers by Client Type
    DHCP                     107
    PPP                      76
    VLAN                      8

    TOTAL                191

Subscribers by LS:RI
    default:default      1
    default:ri1              28
    default:ri2              16
    ls1:default             22
    ls1:riA                 38
    ls1:riB                 44
    logsysX:routinstY     42

    TOTAL                    191
```

show subscribers summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
   Active: 3998
   Total: 3998

Subscribers by Client Type
   DHCP: 3998
   Total: 3998

Subscribers by LS:RI
   default:default: 3998
   Total: 3998
```

show subscribers summary physical-interface pic

```
user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
   Active: 4825
   Total: 4825
```

```
Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
Interface           Count
ge-1/0              1000
ge-1/3              1000

Total Subscribers: 2000
```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
Interface           Count
ae0: ge-1/0         801
ae0: ge-1/3         801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
Interface           Count
ge-5/0/1              201
ge-5/0/2              301

Total Subscribers: 502
```

show subscribers summary port (Pseudowire Interfaces)

```
user@host> show subscribers summary port
ps0: lt-2/1/0 10
ps1: lt-2/1/0 20

Total Subscribers: 30
```

show subscribers summary port extensive

```
user@host>show subscribers summary port extensive
Interface: xe-3/0/3
Port Count: 100
Detail:
Subscribers by Client Type
   PPPoE: 1
   ESSM: 99
Subscribers by Connection Type
   Terminated: 1

Interface: xe-3/1/3
Port Count: 3100
Detail:
Subscribers by Client Type
   PPPoE: 1600
   ESSM: 1100
   VLAN-OOB: 400
Subscribers by Connection Type
   Tunneled: 500
   Terminated: 1100
   Cross-connected: 400

Total Subscribers: 26197
```

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface                      IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824         100                                    default:default
```

```
demux0.1073741824           203.0.113.10         WHOLESALER-CLIENT default:default
demux0.1073741825           203.0.113.13         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826           203.0.113.213        RETAILER2-CLIENT  test1:retailer2
```

show subscribers summary (agf)

```
user@host> show subscribers summary
Subscribers by Sub-System
   AGF: 3000
   BNG: 0
   Total: 3000

Subscribers by State
   Active: 4000
   Total: 4000

Subscribers by Client Type
   VLAN: 1000
   PPPoE: 3000
   Total: 4000
```

show subscribers summary sub-system agf

```
user@host> show subscribers summary sub-system agf
AGF Subscribers by State
    Init:     3
    Active:  12
    Total:   15

AGF Subscribers by Client Type
     DHCP:           2
     PPPOE:          5
     PPPOE-5G        5
     TOTAL          12
```

Release Information

Command introduced in Junos OS Release 10.2.

`sub-system` added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

show subscribers | 3057

show system subscriber-management redundancy-state dhcp active-leasequery interface

IN THIS SECTION

- Syntax | 3121
- Description | 3121
- Options | 3122
- Required Privilege Level | 3122
- Output Fields | 3122
- Sample Output | 3123
- Release Information | 3123

Syntax

```
show system subscriber-management redundancy-state dhcp active-leasequery interface interface-name
```

Description

Display the redundancy state of the specified access logical interface based on the redundancy primary role election protocol, VRRP or pseudowire.

For VRRP redundancy, the redundancy state is the same as the VRRP state of the underlying logical interface:

- The redundancy state is `Master` if the logical interface is in VRRP active mode.

- The redundancy state is `Backup` if the logical interface is in VRRP backup mode.

For pseudowire redundancy, the redundancy state is based on the state of the pseudowire interface:

- The redundancy state is `Master` if the logical interface is in the UP state for a pseudowire interface.
- The redundancy state is `Backup` if the logical interface is in the DOWN state for a pseudowire interface.

Options

interface-name	Name of the access logical interface.

Required Privilege Level

view

Output Fields

Table 160 on page 3122 lists the output fields for the `show system subscriber-management redundancy-state dhcp active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 160: show system subscriber-management redundancy-state dhcp active-leasequery Output Fields

Field Name	Field Description
interface	Name of the access logical interface.
Redundancy State	Redundancy state of the specified interface. • `Master`—Interface is on the current primary BNG. This BNG is handling traffic for all members of the subscriber redundancy group that is on this interface. For VRRP, the interface state is `Master`; for pseudowire, the interface state is `Up`. • Backup—Interface is on the current backup BNG. If the primary BNG fails, this BNG is elected as the new primary. This interface transitions to the `Master` state and handles traffic for all members of the subscriber redundancy group previously handled on the corresponding access interface on the former primary. For VRRP, the interface state is backup; for pseudowire, the interface state is down.

Sample Output

show system subscriber-management redundancy-state dhcp active-leasequery interface (VLAN Underlying Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
ge-0/0/0.1

Interface                    Redundancy State
ge-0/0/0.1                   Master
```

show system subscriber-management redundancy-state dhcp active-leasequery interface (Subscriber Demux Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
demux0.3221225473

Interface                    Redundancy State
demux0.3221225473            Master
```

show system subscriber-management redundancy-state dhcp active-leasequery interface (Pseudowire Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
ps1.0

Interface                    Redundancy State
ps1.0                        Master
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

M:N Subscriber Redundancy on BNG | 1173

show system subscriber-management redundancy-state interface

IN THIS SECTION

- Syntax | 3124
- Description | 3124
- Options | 3124
- Required Privilege Level | 3124
- Output Fields | 3124
- Sample Output | 3125
- Release Information | 3127

Syntax

```
show system subscriber-management redundancy-state interface interface-name
```

Description

Displays the status of the subscriber management redundancy service of the pseudowire interface.

Options

interface *interface-name*	Name of the pseudowire interface.

Required Privilege Level

view

Output Fields

Table 161 on page 3125 lists the output fields for the `show system subscriber-management redundancy-state interface interface-name` command.

Table 161: show system subscriber-management redundancy-state interface Output Fields

Field Name	Field Description
Forwarding state	Status of the packet forwarding.
Service Activation Programming	Status of the service activation programming.
Standby-mode	Status of the standby mode.

Sample Output

Primary BNG Interface Status

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status in Normal Operation

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Secondary BNG Interface Status Immediately on Failover

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: In-progress
```

Secondary BNG Interface Status on Failover Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status Immediately on Failover Reversal

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: In-rollback
```

Secondary BNG Interface Status Immediately on Failover Reversal Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

N+1 Support for BNG M:N Subscriber Service Redundancy | 1223

show system subscriber-management route

IN THIS SECTION

- Syntax | 3128
- Description | 3128
- Options | 3128
- Required Privilege Level | 3129
- Output Fields | 3129
- Sample Output | 3131
- Release Information | 3134

Syntax

```
show system subscriber-management route
<family family>
<incomplete>
<level (brief | detail)>
<next-hop index>
<prefix>
<routing-instance name>>
<route-type type>
<rrt-index index>
<summary>
```

Description

Display information about how routes are mapped to specific enhanced subscriber management interfaces. You can customize the output by including one or more optional filters in the command. With the exception of the `summary` option, all filter options can be combined in a single command.

Options

family *family* (Optional) Display route mapping information for the specified protocol family: `inet` (IPv4) or `inet6` (IPv6).

incomplete (Optional) Display route mapping information for incomplete routes that are missing elements required to add the routes to the routing table.

level (brief | detail) (Optional) Display the specified level of output: `brief` or `detail`.

next-hop *index* (Optional) Display the next hop associated with the route entry with the specified next-hop index, in the range 1 through 65535.

prefix *address* (Optional) Use the same prefix and prefix length as the subscriber host address. Output includes attributes that originate in the Famed-Route record of an upstream RADIUS server (Tag, Metric, Preference).

route-type *type* (Optional) Display route mapping information for the specified route type: `access`, `access-internal`, `kernel`, or `local`.

routing-instance *name* (Optional) Display route mapping information for the specified routing-instance

rrt-index *index* (Optional) Display mapping information for the specified routing table index, in the range 0 through 65535. An `rtt-index` value of 0 (zero) denotes routes in the default routing table managed by enhanced subscriber management.

summary (Optional) Display summary information about the routes managed by enhanced subscriber management.

Required Privilege Level

view

Output Fields

Table 162 on page 3129 lists the output fields for the `show system subscriber-management route` command. Output fields are listed in the approximate order in which they appear.

Table 162: show system subscriber-management route Output Fields

Field Name	Field Description	Level of Output
address	IPv4 or IPv6 address associated with the route entry.	All levels
`Route Type`	One of the following route types: • `Access` • `Access-internal` • `Framed` • `Kernel` • `Local`	All levels
`Interface`	Name of the enhanced subscriber management interface associated with the route entry.	All levels
`Next-hop`	Next-hop associated with the route entry.	All levels

Table 162: show system subscriber-management route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Tag	Reflects the Tag attribute used in the RADIUS Framed-Route type record.	All levels
Metric	Reflects the Metric attribute used in the RADIUS Framed-Route type record.	All levels
Preference	Reflects the Preference attribute used in the RADIUS Framed-Route type record.	All levels
Rtt-index	Value of the routing table index. A value of 0 (zero) denotes a route in the default routing table managed by enhanced subscriber management.	detail
Bbe index	Value of the interface index for the control plane.	detail
Flow id	Value of the route object index.	detail
Reference Count	Used for internal accounting.	detail
Dirty Flags	Used for internal accounting.	detail
Flags	Used for internal accounting.	detail
Family	One of the following protocol families: • AF_INET—IPv4 • AF_INET6—IPv6	detail

Sample Output

show system subscriber-management route prefix <address>

`rtt-index 0`

```
user@host> show system subscriber-management route prefix 10.10.0.1/32
Route:  10.10.0.1/32
     Routing-instance:         default:default
     Kernel rt-table id :      0
     Family:                   AF_INET
     Route Type:               Framed
     Protocol Type:            Unspecified
     Interface:                pp0.3221225491
     Interface index:          26
     Internal Interface index: 26
     Route index:              20
     Next-Hop:                 684
     Tag:                      9999
     Metric:                   56
     Preference:               10
     Reference-count:          1
     L2 Address:               00:00:5e:00:53:0b
     Flags:                    0x0
     Dirty Flags:              0x0
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
```

```
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

show system subscriber-management route family route-type rtt-index level detail

The following example displays detailed information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level detail
2001:db8::/64
     Route Type:      Access
     Interface:       pp0.3221225479
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:1::/64
     Route Type:      Access
     Interface:       pp0.3221225477
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
```

```
     Flow id:          1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:2::/64
     Route Type:      Access
     Interface:       pp0.3221225478
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:3::/64
     Route Type:      Access
     Interface:       pp0.3221225480
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
```

```
2001:db8:0:0:2::/64
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Support for passing **Framed-Route** attributes from a RADIUS server to the router was added in Junos OS Release 17.2 on MX Series routers for enhanced subscriber management. This allows the tagged subscriber host routes to be imported to the routing table and advertised by BGP.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show system subscriber-management statistics

IN THIS SECTION

- Syntax | **3135**
- Description | **3135**

- Options | 3135
- Required Privilege Level | 3136
- Output Fields | 3136
- Sample Output | 3137
- Release Information | 3146

Syntax

```
show system subscriber-management statistics
<all>
<dhcp>
<dvlan>
<fixed-wireless-access>
<l2tp>
<ppp>
<pppoe>
```

Description

Display statistics for the specified option. You can customize the output by including one or more optional filters in the command. With the exception of the `extensive` option, all filter options can be combined in a single command.

Options

all	(Optional) Display packet statistics for all protocols.
dhcp	(Optional) Display DHCP packet statistics.
dvlan	(Optional) Display DVLAN packet statistics.
fixed-wireless-access	(Optional) Display fixed wireless access packet statistics.
l2tp	(Optional) Display L2TP packet statistics.
ppp	(Optional) Display PPP packet statistics.

pppoe	(Optional) Display PPPoE packet statistics.

Required Privilege Level

view

Output Fields

Table 163 on page 3136 lists the output fields for the `show system subscriber-management statistics` command. Output fields are listed in the approximate order in which they appear.

Table 163: show system subscriber-management statistics Output Fields

Field Name	Field Description
`Rx Statistics`	Statistics for packets received.
`Tx Statistics`	Statistics for packets sent.
`Enhanced I/O Statistics`	Statistics for visibility into packet drops from the queue.
`Error Statistics`	Includes connection packets, flow control, and messages and packets sent to and received from the daemon.
`ERA discards`	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
`Layer 3 Statistics`	Statistics for Layer 3 packets.
`padis`	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
`padrs`	PPPoE Active Discovery Request packets.
`ppp`	Point-to-Point Protocol packets.

Table 163: show system subscriber-management statistics Output Fields *(Continued)*

Field Name	Field Description
router solicitations	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.
router advertisements	Number of router advertisements sent or received.
route solicit response packet	Number of router solicitation responses sent or received.

Sample Output

The following examples displays packet statistics accumulated for DHCP, hybrid access, and PPPoE since the last time the session manager was cleared.

show system subscriber-management statistics all

```
user@host> show system subscriber-management statistics all
  user@host> show system subscriber-management statistics all
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660

DHCP Statistics:
```

```
------------------------------------------------------------
    Rx Statistics
        packets                         : 320008
        ERA discards                    : 6274
    Tx Statistics
        transmit request packets        : 320482
        sent packets                    : 320482
    Error Statistics
    Connection Statistics
        no connection packets           : 0

PPPoE Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 486165
        padis                           : 36768
        padrs                           : 35421
        ppp packets                     : 341787
        ERA discards                    : 8249
    Tx Statistics
        packets                         : 70842
        send failures                   : 6240
```

show system subscriber-management statistics dhcp

```
user@host> show system subscriber-management statistics dhcp
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
```

```
        packets                         : 6604660

DHCP Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                         : 320008
        ERA discards                    : 6274
    Tx Statistics
        transmit request packets        : 320482
        sent packets                    : 320482
    Error Statistics
    Connection Statistics
        no connection packets           : 0
```

show system subscriber-management statistics dhcp extensive

```
user@host> show system subscriber-management statistics dhcp extensive
Session Manager started @ Mon Dec  6 06:14:27 2021
Session Manager cleared @ Mon Dec  6 06:14:27 2021
---------------------------------------------------------------
                    Packet Statistics
---------------------------------------------------------------
I/O Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                         : 7986
    Tx Statistics
        packets                         : 15
          l2 inject                     : 15
          l3 inject                     : 0
    Buffer Statistics
        allocations                     : 7990
        frees                           : 7990
        allocation failures             : 0
  Layer 3 Statistics
    Rx Statistics
        packets                         : 29
    Tx Statistics
        packets                         : 15
    PFE Event Statistics
        packets                         : 7957
```

```
-------------------------------------------------------------
Enhanced I/O Statistics:
-------------------------------------------------------------
    bbe_io_rcv l2                        : 0
    bbe_io_rcv l3                        : 29
    bbe_io_rcv l3 v4                     : 17
    bbe_io_rcv l3 v6                     : 12
    bbe_io_rcv l3 unspec                 : 0
    bbe_io_rcv l3 unknown af             : 0
    bbe_io_rcv routed                    : 0
    bbe_io_rcv routed_v4                 : 0
    bbe_io_rcv routed_v6                 : 0
    bbe_io_rcv routed_no_route           : 0
    bbe_io_rcv routed_default            : 0
    bbe_io_rcv resolve v4                : 0
    bbe_io_rcv resolve v6                : 0
    bbe_io_rcv resolve default           : 0
    bbe_io_rcv pfe event                 : 0
    bbe_io_rcv default                   : 0
    -----------------------------------------------------------
    rx inet                              : 17
    rx inet6                             : 12
    rx inet6 icmp6                       : 2
    rx inet igmp                         : 0
    rx inet gre                          : 0
    rx inet agf                          : 0
    rx inet icmp                         : 0
    rx inet udp v4                       : 17
    rx inet udp v6                       : 10
          rx inet udp l2tp               : 0
          rx inet udp jdhcp v4           : 17
          rx inet udp jdhcp v6           : 10
          rx inet udp bfd v4             : 0
          rx inet udp bfd v6             : 0
    rx inet tcp proxy v4                 : 0
    rx inet tcp proxy v6                 : 0
    rx v4 ip frag reassm proc cnt        : 0
    rx v4 ip frag reassm pkt cnt         : 0
    rx v4 ip frag reassm alloc cntx cnt  : 0
    rx v4 ip frag reassm free cntx cnt   : 0
    tx l3 forward                        : 15
    tx udp v4                            : 5
    tx udp v6                            : 6
```

```
    tx tcp proxy client v4               : 0
    tx tcp proxy client v4 error         : 0
    tx tcp proxy v4 drop                 : 0
    tx tcp proxy v6 drop                 : 0
    tx igmp                              : 0
    tx v6 l3 route forwards              : 0
    tx v6 l3 route drops                 : 0
    tx v6 kernel forwards                : 0
    tx v6 kernel forward drops           : 0
    tx v6 l2 forwards                    : 10
    tx v6 l2 drops                       : 0
    tx v4 l3 route forwards              : 0
    tx v4 l3 route drops                 : 0
    tx v4 kernel forwards                : 0
    tx v4 kernel forward drops           : 0
    tx v4 l2 forwards                    : 5
    tx v4 l2 drops                       : 0
    tx v4 ip fragment forward            : 0
    tx v4 ip fragment DF drops           : 0
    tx v4 ip fragment drops              : 0
    tx v4 ip fragment failed             : 0
    bbe_ifl_output tx failed             : 0
    bbe_io_send tx failed                : 0
    bbe_io_send tx partial failed        : 0
    io_queue low_prio_packets_dropped    : 0
    io_queue mlow_prio_packets_dropped   : 0
    io_queue med_prio_packets_dropped    : 0
    io_queue high_prio_packets_dropped   : 0
-------------------------------------------------------------
DHCP Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 27
    Tx Statistics
        transmit request packets         : 11
        sent packets                     : 11
    DHCPv4 Rx Statistics
        total packets                    : 17
        boot request                     : 15
        boot reply                       : 2
        discover                         : 11
        offer                            : 1
        request                          : 3
```

```
        ack                             : 1
        release                         : 1
    DHCPv4 Tx Statistics
        total packets                   : 5
        boot reply                      : 5
        offer                           : 2
        ack                             : 3
    DHCPv6 Rx Statistics
        total packets                   : 10
        solicit                         : 4
        advertise                       : 1
        request                         : 2
        reply                           : 1
        renew                           : 1
        release                         : 1
        relay repl                      : 2
    DHCPv6 Tx Statistics
        total packets                   : 6
        advertise                       : 2
        reply                           : 4
    Error Statistics
    Connection Statistics
        no connection packets           : 0
        connection down events          : 1
        connection up events            : 2
        flow control invoked            : 0
        flow control released           : 0
        packets sent to daemon          : 27
        packets received from daemon    : 11
        messages sent to daemon         : 0
        messages received from daemon   : 1481
        notifies while not connected    : 0
NET Statistics:
------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                        : 2
        neighbor advertisements         : 2
    Tx Statistics
        packets:                        : 4
        neighbor solicitations          : 4
Management Statistics:
------------------------------------------------------------
```

```
    ipdemux                             : 4
    ipdemux add                         : 4
    ipdemux delete                      : 2
    ifl flow                            : 7
    ifl flow adds                       : 6
    ifl flow changes                    : 1
    ifl flow deletes                    : 3
    ip flow                             : 9
    ip flow add                         : 6
    ip flow delete                      : 3
    service                             : 27
    service adds                        : 24
    service deletes                     : 3
Management Config Status:
------------------------------------------------------------
    gres state enabled state            : 0
    shmlog disabled state               : 0
    shmlog filtering state              : 0
    vc backup member local switch state  : 0
    dscp code point value               : 0x30
```

show system subscriber-management statistics pppoe

```
user@host> show system subscriber-management statistics pppoe
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660
```

```
PPPoE Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 486165
        padis                           : 36768
        padrs                           : 35421
        ppp packets                     : 341787
        ERA discards                    : 8249
    Tx Statistics
        packets                         : 70842
        send failures                   : 6240
```

show system subscriber-management statistics extensive

```
user@host> show system subscriber-management statistics extensive
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
    Buffer Statistics
        allocations                     : 7032618
        frees                           : 7032624
        allocation failures             : 0
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660
    PFE Event Statistics
        packets                         : 0

------------------------------------------------------------
Enhanced I/O Statistics:
------------------------------------------------------------
```

```
    bbe_io_rcv l2                       : 0
    bbe_io_rcv l3                       : 0
    bbe_io_rcv l3 v4                    : 0

    io low queue drops                  :12
    io mlow queue drops                 :0
    io medium queue drops               :0
    io high queue drops                 :0
```

show system subscriber-management statistics ppp (LCP Vendor-Specific Counters)

```
user@host> show system subscriber-management statistics ppp
Session Manager started @ Thu Feb 11 00:37:43 2020
Session Manager cleared @ Thu Feb 11 00:37:43 2020
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 486783
    Tx Statistics
        packets                         : 144
  Layer 3 Statistics
    Rx Statistics
        packets                         : 8
    Tx Statistics
        packets                         : 0
PPP Statistics:
------------------------------------------------------------
    Rx Statistics
        network packets                 : 123
        plugin packets                  : 123
        lcp config requests             : 18
        lcp config acks                 : 18
        lcp conf nacks                  : 8
        lcp conf rejects                : 6
        lcp termination requests        : 4
        lcp termination acks            : 13
        lcp code rejects                   : 2
        lcp vendor-specific acks       : 10
```

```
        pap requests                    : 8
        ipcp requests                   : 27
        ipcp acks                       : 9
        ipv6cp requests                 : 11
        ipv6cp acks                     : 1
    Tx Statistics
        packets                         : 101
        lcp config requests             : 32
        lcp config acks                 : 18
        lcp termination requests        : 13
        lcp termination acks            : 4
        lcp vendor-specific requests    : 10
        pap acks                        : 8
        ipcp requests                   : 9
        ipcp acks                       : 5
        ipcp nacks                      : 9
        ipv6cp requests                 : 1
        ipv6cp acks                     : 1
        ipv6cp nacks                    : 1
NET Statistics:
------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                        : 8
        router solicitations            : 8
    Tx Statistics
        packets:                        : 0
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Enhanced I/O Statistics introduced as part of Extensive output in Junos OS Release 15.1R4 on MX Series routers for enhanced subscriber management.

RELATED DOCUMENTATION

- Understanding Dropped Packets and Untransmitted Traffic Using show Commands

show system subscriber-management summary

IN THIS SECTION

- Syntax | 3147
- Description | 3147
- Options | 3147
- Required Privilege Level | 3147
- Output Fields | 3147
- Sample Output | 3150
- Release Information | 3151

Syntax

```
show system subscriber-management summary
```

Description

Display complete subscriber management database summary information.

Options

none	This command has no options.

Required Privilege Level

view

Output Fields

Table 164 on page 3148 lists the output fields for the `show system subscriber-management summary` command. Output fields are listed in the approximate order in which they appear.

Table 164: show system subscriber-management summary Output Fields

Field Name	Field Description
`Graceful Restart`	State of graceful Routing Engine switchover (GRES): • `Enabled` • `Disabled` (Enhanced subscriber management for MX Series routers) The name of this field is `Graceful Switchover`.
`Mastership`	State of the Routing Engine: • `Master` • `Standby`
`Database`	State of the subscriber management database: • `Available` • `Init` • `Not-available`
`Standby`	(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine: • `Connected`—Connected but not synchronized • `Disconnected`—Not connected • `Resync (`*nn*`%)`—Connected and *nn* percent synchronized with the primary Routing Engine • `Synchronized`—Synchronized with the primary Routing Engine
`Disconnection Reason`	Reason why both Routing Engines are disconnected when there is a DRAM mismatch. • Primary/Standy RE DRAM Size Mismatch—Displayed when the amount of memory is different on the primary and standby Routing Engines.

Table 164: show system subscriber-management summary Output Fields *(Continued)*

Field Name	Field Description
Chassisd ISSU State	State of unified ISSU chassis daemon: • ABORT • DAEMON_ISSU_PREPARE • DAEMON_ISSU_PREPARE_DONE • DAEMON_SWITCHOVER_PREPARE • DAEMON_SWITCHOVER_PREPARE_DONE • FRU_ISSU • FRU_ISSU_DONE • IDLE • UNKNOWN
ISSU State	State of unified ISSU aggregate daemon: • ABORT • IDLE • PREPARE • READY • SWITCHOVER_PREPARE • SWITCHOVER_READY • UNKNOWN
ISSU Wait	Amount of time, in seconds, requested by a daemon to perform cleanup. If multiple daemons request time, the displayed value is the highest wait time requested by a daemon.

Sample Output

show system subscriber-management summary

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

show system subscriber-management summary (Enhanced Subscriber Management)

```
user@host> show system subscriber-management summary
General:
    Graceful Switchover     Enabled
    Mastership              Master
    Database                Available
    Standby                   Resync (75%)
    Chassisd ISSU State     IDLE
    ISSU State              IDLE
    ISSU Wait               0
```

show system subscriber-management summary (DRAM Size Mismatch Error)

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Standby              Disconnected
<emphasis>    Disconnection Reason    Master/Standy RE DRAM Size Mismatch</emphasis
>
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show database-replication statistics | 2776

show database-replication summary | 2778

test aaa authd-lite user

IN THIS SECTION

- Syntax | 3151
- Description | 3151
- Options | 3152
- Required Privilege Level | 3152
- Output Fields | 3152
- Sample Output | 3153
- Release Information | 3156

Syntax

```
test aaa authd-lite user username password password profile access-profile-name
<port nas-port>
<zero-stats>
```

Description

Verify authd-lite subscriber access authentication, accounting, and address allocation configuration.

The `test aaa` command supports 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.

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. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username	Specify the subscriber username to test.
password *password*	Specify the password associated with the username.
profile *access-profile-name*	Specify the access profile associated with the subscriber.
port *nas-port*	(Optional) Specify the NAS port used for the test.
zero-stats	(Optional) Specify that no accounting statistics are set for this test.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa authd-lite user

The following example tests the configuration for authd-lite subscriber user1bt with a password of $ABC123 and an access profile of employee12, and displays the resulting output:

```
user@host> test aaa authd-lite user user1bt password $ABC123 profile employee12
    Authentication Grant
    ************User Attributes************
         User Name -                          user1bt
         Framed IPv6 Prefix -                 ::/0
         Framed IPv6 Pool -                   NULL
         Nas IPv6 Address -                   ::
         NDRA IPv6 Prefix -                   NULL
         Login IPv6 Host -                    ::
         Framed Interface Id -                0:0:0:0
         Delegated IPv6 Prefix -              ::/0
         NDRA IPv6 Pool -                     NULL
         User Password -                      $ABC123
         Nas Ip Address -                     0.0.0.0
         NAS Port -                           0
         Service Type-                        0
         Framed IP Address -                  0.0.0.0
         Framed IP Netmask -                  0.0.0.0
         Filter Id -                          NULL
         Framed MTU -                         0
         Reply Message -                      NULL
         Framed Route-                       <not set>
         Framed MTU -                         0
         Class -                              SBR2CL
         Virtual Router Name (LS:RI)          default:default
         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
         Ingress Statistics -                 disabled
         Egress Statistics -                  disabled
         Ingress Policy Name -               <not set>
         Engress Policy Name -               <not set>
         IGMP Enable -                        disabled
         Redirect VR Name (LS:RI)             default:default
```

```
      Service Bundle                          <not set>
      Framed Ip Route Tag                     <not set>
      LI Action                               0
      LI Interception Identifier              0
      LI Mediation Device IP Address          0.0.0.0
      LI_Mediation_Device_Port_Number         0
      Activate Service                        NULL
      Deactivate Service                      NULL
      Service Statistics                      0
      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>
      MLD Version -                           MLD Version not set
      IGMP Version                            IGMP Version not set
      IGMP Immediate Leave -                  <not set>
      MLD Immediate Leave -                   <not set>
      IPv6_Ingress_Policy_Name -              <not set>
      IPv6_Egress_Policy_Name -               <not set>
      Cos_Parameter_Type -                    <not set>
      Service Interim Acct Interval           0
      Max Clients Per Interface               0
      Cos_Scheduler_Pmt_Type                  <not set>
      Session Timeout                         599999940
      NAS Port Type                           0
      Framed Pool                             NULL
      Idle Timeout                            0
Acct-start sent
Acct-start succeeded
Pausing 10 seconds
Interim-Acct sent
Acct-interim succeeded
Pausing 10 seconds
Acct-stop sent
Acct-stop succeeded
Logging out subscriber
Test complete. Exiting
```

test aaa authd-lite user (XML Output, Old Format)

The following example shows an excerpt of sample XML output in the old format:

```
user@host>test aaa authd-lite user user45@test.net password $ABC123 profile test | 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>Framed IPv6 Prefix -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        <radius-server-attribute-name>Framed IPv6 Pool -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        <radius-server-attribute-name>NDRA IPv6 Prefix -</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>
```

test aaa authd-lite user (XML Output, New Format)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa authd-lite 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>Framed IPv6 Prefix -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Framed IPv6 Pool -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>NDRA IPv6 Prefix -</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>
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Testing a Subscriber AAA Configuration | 676

Understanding External Authentication Servers

test aaa dhcp user

IN THIS SECTION

- Syntax | 3157
- Description | 3157
- Options | 3158

- Required Privilege Level | **3159**
- Output Fields | **3159**
- Sample Output | **3160**
- Release Information | **3163**

Syntax

```
test aaa dhcp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<mac-address mac-address>
<no-address-request>
<option-82 option-82>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
<service-type service-type>
<source-address source-address>
<terminate-code code-value>
```

Description

Verify Dynamic Host Configuration Protocol (DHCP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

NOTE: The `test aaa` command supports 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: 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. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username	Subscriber username to test.
agent-remote-id *ari*	(Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).
logical-system *logical-system-name*	(Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).
mac-address *mac-address*	(Optional) MAC address of the DHCP client.
no-address-request	(Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed. **NOTE**: The `test aaa dhcp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.
option-82 *option-82*	(Optional) DHCP relay agent information option (option-82) value.
password *password*	(Optional) Password associated with the username.
profile *access-profile-name*	(Optional) Access profile associated with the subscriber.
routing-instance *routing-instance-name*	(Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.
service-type *service-type*	(Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings

that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)
callback-login (3)	–

source-address *source-address* (Optional) IP address of the outgoing interface.

terminate-code *code-value* (Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa dhcp user

The following example tests the configuration for DHCP subscriber user1DB and password $ABC123, and displays the resulting output:

```
user@host> test aaa dhcp user user1DB@test.net password $ABC123
Authentication Grant
    ************User Attributes************
         User Name -                          user1DB@test.net
         Client IP Address -                  192.168.1.1
         Client IP Netmask -                  255.255.0.0
         Virtual Router Name (LS:RI)-         default: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>
         Service Type -                       0
         DHCP Guided Relay Server -           0
         Class Attribute -                    TEST
         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 Enable -                           disabled
        Redirect VR Name (LS:RI)-               default: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 -                 <not set>
        MLD Immediate Leave -                  <not set>
        IPv6 Ingress Policy Name -             <not set>
        IPv6 Egress Policy Name -              <not set>
        Dynamic Profile -                      <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
```

test aaa dhcp user (XML Output, Old Format)

The following example shows an excerpt of sample XML output in the old format:

```
user@host>test aaa dhcp user user45@test.net password $ABC123 profile test | 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>Client IP Address -</radius-server-attribute-name>
        <radius-server-attribute-value>198.51.100.7</radius-server-attribute-value>
        <radius-server-attribute-name>Client IP Netmask -</radius-server-attribute-name>
        <radius-server-attribute-value>255.255.255.255</radius-server-attribute-value>

...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa dhcp user (XML Output, New Format)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa dhcp 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>Client IP Address -</radius-server-attribute-name>
            <radius-server-attribute-value>198.51.100.7</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Client IP Netmask -</radius-server-attribute-name>
            <radius-server-attribute-value>255.255.255.255</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>

...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `terminate-code` added in Junos OS Release 11.4.

Option `agent-remote-id` added in Junos OS Release 14.1.

Options `no-address-request` and `service-type` added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Testing a Subscriber AAA Configuration | 676

test aaa ppp user

IN THIS SECTION

- Syntax | 3164
- Description | 3164
- Options | 3165
- Required Privilege Level | 3167
- Output Fields | 3167
- Sample Output | 3167
- Release Information | 3173

Syntax

```
test aaa ppp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<no-address-request>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
<service-type service-type>
<terminate-code code-value>
```

Description

Verify Point-to-Point Protocol (PPP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

NOTE: The `test aaa` command supports 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: 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. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username	Subscriber username to test.
agent-remote-id *ari*	(Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).
logical-system *logical-system-name*	(Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).
no-address-request	(Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed. **NOTE**: The `test aaa ppp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.
password *password*	(Optional) Password associated with the username.
profile *access-profile-name*	(Optional) Access profile associated with the subscriber.

NOTE: The system logically treats this profile as a client-level configuration. An access profile configured in a domain map takes precedence over client-level configurations. If you have configured one or more domain maps, the username for the user under test is evaluated against the domain maps the same as any other subscriber.

For example, the username can exactly match a domain map or partially match a wildcard domain map. If it matches neither of those, then it matches the `default` domain map if it is configured. If the username has no domain or realm ,then it matches the `none` domain map, if it is configured.

The consequence is that if the test user matches any configured domain map, then an access profile configured in that map is used for the test in preference to an access profile that you specify with the `test` command.

See Specifying an Access Profile in a Domain Map for more information about domain maps and access profiles.

routing-instance *routing-instance-name*

(Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.

service-type *service-type*

(Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)

callback-login (3)	–

terminate-code *code-value* (Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is <not set>.

Sample Output

test aaa ppp user

The following example tests the configuration for PPP subscriber user98BEDC and password $ABC123, and displays the resulting output:

```
user@host> test aaa ppp user user98BEDC@test.net password $ABC123
Authentication Grant
    ************User Attributes***********
         User Name -                              user98BEDC@test.net
         Client IP Address -                      192.168.1.1
         Client IP Netmask -                      255.255.0.0
         Virtual Router Name (LS:RI) -            default: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 Enable -                           disabled
        Redirect VR Name (LS:RI) -              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 -                  <not set>
          MLD Immediate Leave -                   <not set>
          IPv6 Ingress Policy Name -              <not set>
          IPv6 Egress Policy Name -               <not set>
          Dynamic Profile -                       <not set>
          Acct Session ID -                        1
          Acct Interim Interval -                  750
          Acct Type -                              1
          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
```

test aaa ppp user (tunneled user)

The following example tests the configuration for PPP tunneled subscriber accounting14, with password $ABC123 and access profile finance-b, and displays the resulting output:

```
user@host> test aaa ppp user accounting14 password $ABC123 14 profile finance-b
    Authentication Grant with Tunnel Attributes
    ************Tunnel Attributes***********
         ****Tunnel Definiton -                  1
             Tunnel Medium           -           1
             Tunnel Type             -           3
             Tunnel Max Sessions     -           100
             Tunnel Server Endpoint  -           192.0.2.4
             Tunnel Client Endpoint  -           198.51.100.5
```

```
            Tunnel Server AuthId    -            rt1
            Tunnel Client AuthId    -            ts1
            Tunnel Password         -            radius
            Tunnel Assignment Id    -            til
            Tunnel Logical System   -
            Tunnel Routing Instance -
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                          l2tp session-receive-cdn-avp-bad-hidden
    Test complete. Exiting
```

test aaa ppp user (authentication failure)

The following example shows sample output when the authentication grant fails due to an invalid password:

```
user@host>test aaa ppp user user45@test.net password $ABC123123
 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 (LS:RI)-         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 -                         $ABC123123
          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 Enable-                            disabled
          Redirect VR Name (LS:RI)-               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 -                  not set
          MLD Immediate Leave -                   not set
          IPv6 Ingress Policy Name -              not set
          IPv6 Egress Policy Name -               not set
          Acct Session ID -                       12
          Acct Interim Interval -                 0
          Acct Type -                             0                       Chargeable user
identity -              0
          NAS Port Id -                           -0/0/0.0
          NAS Port -                              4095
          NAS Port Type -                         15
          Framed Protocol -                       0
    Test complete. Exiting
```

test aaa ppp user (XML Output, Old 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>
```

test aaa ppp user (XML Output, New Format)

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>
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `terminate-code` added in Junos OS Release 11.4.

Option `agent-remote-id` added in Junos OS Release 14.1.

Options `no-address-request` and `service-type` added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Testing a Subscriber AAA Configuration | 676