Junos® OS

Broadband Subscriber Sessions Feature Guide

Modified: 2017-05-26

Copyright © 2017, Juniper Networks, Inc.

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Junos® OS Broadband Subscriber Sessions Feature Guide
Copyright © 2017, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation xlv
Documentation and Release Notes xlv
Supported Platforms xlv
Using the Examples in This Manual xlv
Merging a Full Example xlvi
Merging a Snippet xlvi
Documentation Conventions xlvii
Documentation Feedback xlix
Requesting Technical Support xlix
Self-Help Online Tools and Resources xlix
Opening a Case with JTAC l

Part 1 Configuring AAA for Subscriber Management

Chapter 1 AAA and RADIUS for Subscriber Access Overview 3

AAA Service Framework Overview 3
RADIUS Server Options for Subscriber Access 4
Global RADIUS Options for Subscriber Access 8
Subscriber Access Interface Description Storage and Reporting Through RADIUS
Overview 8
Interface Description Precedence 9
Example: Reporting Interface Descriptions on Non-Underlying Logical
Interfaces 9
Reporting Interface Descriptions on Underlying Logical Interfaces 10
Interface Descriptions on Aggregated Ethernet Physical Interfaces 11
Interface Descriptions on a Combination of Dynamic and Static
Interfaces 11
Example: Reporting Interface Descriptions on Dynamic VLANs 11

Chapter 2 Configuring RADIUS Attributes and Juniper Networks VSAs 13

RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service
Framework 13
RADIUS IETF Attributes Supported by the AAA Service Framework 14
Juniper Networks VSAs Supported by the AAA Service Framework 21
AAA Access Messages and Supported RADIUS Attributes and Juniper Networks
VSAs for Junos OS 33
AAA Accounting Messages and Supported RADIUS Attributes and Juniper
Networks VSAs for Junos OS 39
Configuring How RADIUS Attributes Are Used for Subscriber Access 43
Junos OS Predefined Variables That Correspond to RADIUS Attributes and
VSAs 49

Copyright © 2017, Juniper Networks, Inc.

DSL Forum Vendor-Specific Attributes 56
DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS 58
RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses 59

Chapter 3 Configuring RADIUS NAS-Port Attributes and Options 61

Manual Configuration of the NAS-Port-ID RADIUS Attribute 61
Configuring a NAS-Port-ID with Additional Options 63
Configuring the Order in Which Optional Values Appear in the NAS-Port-ID 64
Configuring a Calling-Station-ID with Additional Attributes 66
Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers 70
RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview 71
NAS-Port-Type RADIUS Attribute 72
NAS-Port RADIUS Attribute 72
NAS-Port Options Configuration and Subscriber Network Access Models 72
NAS-Port Options Definition 72
Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN 73
Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN 74
Manual Configuration of the NAS-Port-Type RADIUS Attribute 75
Configuring the RADIUS NAS-Port-Type per Physical Interface 77
Configuring the RADIUS NAS-Port-Type per VLAN 79
Configuring the RADIUS NAS-Port-Type per Stacked VLAN 80
Configuring the RADIUS NAS-Port Extended Format per Physical Interface 82
Configuring the RADIUS NAS-Port Extended Format per VLAN 84
Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN 85
Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces 88

Chapter 4 Configuring RADIUS Authentication for Subscriber Access 91

Retaining Authentication and Accounting Information During Session Startup 91
Configuring Authentication and Accounting Parameters for Subscriber Access 92
Specifying the Authentication and Accounting Methods for Subscriber Access 92
Specifying RADIUS Authentication and Accounting Servers for Subscriber Access 93
Example: Configuring RADIUS-Based Subscriber Authentication and Accounting 94

Chapter 5 Configuring RADIUS Accounting for Subscriber Access 97

RADIUS Accounting Statistics for Subscriber Access Overview 97
Understanding RADIUS Accounting Duplicate Reporting 99
Layer 3 Wholesale Scenarios 99
Other Scenarios 100
Filters for Duplicate Accounting Reports 100
Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting 101

Copyright © 2017, Juniper Networks, Inc.

Preservation of RADIUS Accounting Information During an Accounting Server Outage . . . 102
RADIUS Acct-On and Acct-Off Messages . . . 105
Configuring Per-Subscriber Session Accounting . . . 105
Configuring Per-Service Session Accounting . . . 108
Configuring Service Packet Counting . . . 109
Configuring Back-up Options for RADIUS Accounting . . . 111
Configuring Service Interim Accounting . . . 112

Chapter 6 Configuring Routers and RADIUS Servers for Subscriber Access . . . 113

Configuring Router or Switch Interaction with RADIUS Servers . . . 113
Forcing the Router to Contact the Accounting Server Immediately . . . 115
Configuring RADIUS Server Parameters for Subscriber Access . . . 116
Configuring RADIUS Server Options for Subscriber Access . . . 117
Configuring RADIUS Options for Subscriber Access Globally . . . 120
Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable . . . 121

Chapter 7 Configuring Access Profiles for Subscriber Access . . . 123

Configuring an Access Profile for Subscriber Management . . . 123
Attaching Access Profiles . . . 124

Chapter 8 Configuring the Effect of Service Activation Failures on Subscriber Login . . . 125

Subscriber Session Logins and Service Activation Failures Overview . . . 125
Service and Network Family Activation Process . . . 126
Configuring How Service Activation Failures Affect Subscriber Login . . . 130

Chapter 9 Configuring Session Options for Subscriber Access . . . 133

Understanding Session Options for Subscriber Access . . . 133
Subscriber Session Timeouts . . . 133
Subscriber Username Modification . . . 136
Configuring Subscriber Session Timeout Options . . . 138
Removing Inactive Dynamic Subscriber VLANs . . . 139

Chapter 10 Receiving DHCP Options From a RADIUS Server . . . 141

Centrally Configured Opaque DHCP Options . . . 142
Data Flow for RADIUS-Sourced DHCP Options . . . 144
Multiple VSA 26-55 Instances Configuration . . . 145
DHCP Options That Cannot Be Centrally Configured . . . 145
Monitoring DHCP Options Configured on RADIUS Servers . . . 146

Chapter 11 Configuring RADIUS Logical Line Identification . . . 149

RADIUS Logical Line Identifier (LLID) Overview . . . 149
RADIUS Attributes for LLID Preauthentication Requests . . . 150
Configuring Logical Line Identification (LLID) Preauthentication . . . 152
Configuring a Port and Password for LLID Preauthentication Requests . . . 153

Chapter 12 Provisioning Subscriber Services Using Cisco VSAs . . . 155

Processing Cisco VSAs in RADIUS Messages for Service Provisioning . . . 155
Configuring Service Interim Accounting . . . 158

Copyright © 2017, Juniper Networks, Inc.

Chapter 13 Configuring Domain Maps for Subscriber Management 159
Domain Mapping Overview 160
Types of Domain Maps and Their Order of Precedence 161
Wildcard Domain Map 162
Default Domain Map 162
Domain Map for Subscriber Usernames With No Domain or Realm Name 162
Understanding Domain Maps and Logical System/Routing Instance Contexts 163
Configuring a Domain Map 164
Configuring a Wildcard Domain Map 165
Specifying an Access Profile in a Domain Map 167
Specifying an Address Pool in a Domain Map 168
Specifying a Dynamic Profile in a Domain Map 168
Specifying an AAA Logical System/Routing Instance in a Domain Map 169
Specifying a Target Logical System/Routing Instance in a Domain Map 170
Configuring Domain and Realm Name Usage for Domain Maps 171
Specifying Domain and Realm Name Delimiters 172
Specifying the Parsing Order for Domain and Realm Names 173
Specifying the Parsing Direction for Domain and Realm Names 173
Enabling Domain Name Stripping 174
Specifying a Tunnel Profile in a Domain Map 175
Specifying a Tunnel Switch Profile in a Domain Map 175
Configuring PADN Parameters for a Domain Map 176
Chapter 14 Configuring Dynamic Service Activation for Subscriber Access 177
Using RADIUS Dynamic Requests for Subscriber Access Management 177
Dynamic Service Activation During Login Overview 178
Configuring RADIUS-Initiated Dynamic Request Support 178
RADIUS-Initiated Change of Authorization (CoA) Overview 179
CoA Messages 179
Qualifications for Change of Authorization 180
Message Exchange 181
Usage Thresholds for Subscriber Services 182
RADIUS-Initiated Disconnect Overview 183
Disconnect Messages 183
Qualifications for Disconnect 183
Message Exchange 183
Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests 184
Chapter 15 Configuring Termination Causes for Protocols 187
Understanding Session Termination Causes and RADIUS Termination Cause Codes 187
Mapping Session Termination Causes to Custom Termination Cause Codes . . . 189
AAA Termination Causes and Code Values 190
DHCP Termination Causes and Code Values 192
L2TP Termination Causes and Code Values 193
PPP Termination Causes and Code Values 213
VLAN Termination Causes and Code Values 221

Copyright © 2017, Juniper Networks, Inc.

Chapter 16 Configuring Extensible Subscriber Services Manager 225
Extensible Subscriber Services Manager Overview 225
Understanding the Dictionary File 226
Chapter 17 Monitoring and Managing AAA Information for Subscriber Access 227
Verifying and Managing Subscriber AAA Information 227
Monitoring Pending RADIUS Accounting Stop Messages 228
Verifying and Managing the RADIUS Dynamic-Request Feature 229
Verifying and Managing Domain Map Configuration 230
Verifying and Managing LLID Preauthentication Configuration 230
Part 2 Configuring DHCP for Subscriber Management
Chapter 18 Using DHCP Overview 233
Extended DHCP Local Server Overview 234
Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools 236
Providing DHCP Client Configuration Information 237
Minimal Configuration for Clients 238
DHCP Local Server and Address-Assignment Pools 238
Extended DHCP Relay Agent Overview 240
Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers 241
DHCP Liveness Detection 242
DHCP Relay Proxy Overview 243
Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers . . . 243
DHCP Attributes Overview 244
Chapter 19 Configuring Default Services That are Activated at Subscriber Login 247
Default Subscriber Service Overview 247
Configuring a Default Subscriber Service 248
Chapter 20 Assigning IP Addresses 249
DHCP Attributes for Address-Assignment Pools 249
Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use 252
Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option 253
Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address 254
Specifying the Subnet for DHCP Client Address Assignment 255
Example: Extended DHCP Local Server Configuration with Optional Pool Matching 255
DNS Address Assignment Precedence 256
Chapter 21 Configuring Lease Times for IP Addresses 257
Configuring a DHCP Lease-Time Threshold 257
DHCP Lease-Time Validation Overview 259
DHCP Lease Timers 260

Copyright © 2017, Juniper Networks, Inc.

DHCP Asymmetric Leasing Overview 261
Configuring DHCP Asymmetric Leasing 263

Chapter 22 Configuring DHCP Individual and Bulk Leasequery 265

DHCP Leasequery and Bulk Leasequery 265
Guidelines for Configuring Support for Leasequery and Bulk Leasequery Operations 266
Configuring and Using DHCP Individual Leasequery 267
Configuring and Using DHCP Bulk Leasequery 270
Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database 273
Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration 274

Chapter 23 Requesting DHCP Client Configuration Information From an Address Pool 277

DHCP Local Server Handling of Client Information Request Messages 277
Enabling Processing of Client Information Requests 278

Chapter 24 Authenticating DHCP Clients Using An External AAA Authentication Service 281

Using External AAA Authentication Services with DHCP 281
Creating Unique Usernames for DHCP Clients 282
Configuring Passwords for Usernames 285

Chapter 25 Grouping Interfaces and Applying a Common DHCP Configuration to the Group 287

Grouping Interfaces with Common DHCP Configurations 287
Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces . . . 288
Configuring Group-Specific DHCP Local Server Options 289
Configuring Group-Specific DHCP Relay Options 290
Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing 291

Chapter 26 Configuring the Number of DHCP Clients Per Interface 297

Specifying the Maximum Number of DHCP Clients Per Interface 297
Allowing Only One DHCP Client Per Interface 298

Chapter 27 Maintaining Subscribers During Interface Delete Events 301

Subscriber Binding Retention During Interface Delete Events 301
Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events 302
Verifying and Managing the DHCP Maintain Subscribers Feature 302

Chapter 28 Forcing Dynamic Reconfiguration of Clients From a DHCP Local Server 305

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients 305
Default Client/Server Interaction 305
Dynamic Client/Server Interaction for DHCPv4 306
Dynamic Client/Server Interaction for DHCPv6 306

Copyright © 2017, Juniper Networks, Inc.

Manually Forcing the Local Server to Initiate the Reconfiguration Process . . . 307
Action Taken for Events That Occur During a Reconfiguration . . . 307
Configuring Dynamic Client Reconfiguration of Extended Local Server Clients . . . 308
Configuring Dynamic Reconfiguration Attempts for DHCP Clients . . . 310
Configuring Deletion of the Client When Dynamic Reconfiguration Fails . . . 311
Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect . . . 312

Chapter 29 Conserving IP Addresses Using DHCP Auto Logout . . . 313

DHCP Auto Logout Overview . . . 313
Auto Logout Overview . . . 313
How DHCP Identifies and Releases Clients . . . 314
Option 60 and Option 82 Requirements . . . 315
Automatically Logging Out DHCP Clients . . . 315
How DHCP Relay Agent Uses Option 82 for Auto Logout . . . 316
DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers . . . 317
Automatically Logging Out DHCPv6 Clients . . . 318

Chapter 30 Overriding Default DHCP Local Server Configuration Settings . . . 321

Overriding Default DHCP Local Server Configuration Settings . . . 321
Overriding the Default DHCP Relay Configuration Settings . . . 324
Deleting DHCP Local Server and DHCP Relay Override Settings . . . 327
Configuring DHCP Attributes Outside Address Assignment Pools . . . 327

Chapter 31 Configuring DHCP Renegotiation While in Bound State . . . 329

DHCP Behavior When Renegotiating While in Bound State . . . 329

Chapter 32 Attaching Access Profiles to DHCP Interfaces . . . 331

Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview . . . 331
Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces . . . 332
Attaching an Access Profile to All DHCP Subscriber or All DHCP Client Interfaces . . . 332
Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients . . . 332
Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces . . . 333

Chapter 33 Applying a Common DHCP Relay Agent Configuration to Groups of DHCP Servers . . . 335

Configuring Named Server Groups . . . 335
Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups . . . 335

Chapter 34 Changing the Gateway IP Address (giaddr) Field and DHCP Relay Request and Release Packet Source Address . . . 337

Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent . . . 337
Replacing the DHCP Relay Request and Release Packet Source Address . . . 337

Copyright © 2017, Juniper Networks, Inc.

Chapter 35 Securing DHCP Messages Sent Between DHCP Clients and Servers in Different VRFs . . . 339
DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs . . . 339
Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs . . . 340
Chapter 36 Configuring DHCP Relay Agent . . . 345
DHCP-Initiated Service Change Based on Remote ID . . . 345
Using Layer 2 Unicast Transmission for DHCP Packets . . . 346
Trusting Option 82 Information . . . 346
Sending Release Messages When Clients Are Deleted . . . 347
Disabling Automatic Binding of Stray DHCP Requests . . . 348
Using DHCP Relay Agent Option 82 Information . . . 349
Configuring Option 82 Information . . . 350
Including a Prefix in DHCP Options . . . 352
Including a Textual Description in DHCP Options . . . 354
Using DHCP Option Information to Selectively Process DHCP Client Traffic . . . 356
Extracting an Option 82 or Option 37 Substring to Create an Interface Set . . . 358
Configuring DHCP Relay Agent Source Address . . . 360
Example: Minimum DHCP Relay Agent Configuration . . . 360
Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings . . . 361
Example: DHCP Relay Agent Configuration with Multiple Clients and Servers . . . 365
Example: Configuring DHCP-Initiated Service Change Based on Remote ID . . . 366
Disabling DHCP Relay . . . 368
Chapter 37 Configuring DHCP Relay Proxy Mode . . . 369
Enabling DHCP Relay Proxy Mode . . . 369
Chapter 38 Configuring DHCP Local Server Authentication . . . 371
Configuring a Token for DHCP Local Server Authentication . . . 371
Chapter 39 Configuring a Minimum DHCP Local Server Configuration . . . 373
Example: Minimum Extended DHCP Local Server Configuration . . . 373
Chapter 40 Protecting the Routing Engine Using DHCP Firewall Filters . . . 375
Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine . . . 375
Port Number Requirements for DHCP Firewall Filters . . . 379
Chapter 41 Monitoring and Managing DHCP . . . 381
Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings . . . 381
Clearing DHCP Bindings for Subscriber Access . . . 382
Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings . . . 384
Monitoring DHCP Relay Server Responsiveness . . . 384
Verifying and Managing DHCP Local Server Configuration . . . 385
Verifying and Managing DHCP Relay Configuration . . . 385

Copyright © 2017, Juniper Networks, Inc.

Part 3 Configuring IPv6 for Subscriber Management

Chapter 42 Migrating to IPv6 Using IPv4 and IPv6 Dual Stack 389

Why Use IPv4/IPv6 Dual Stack? 389
Basic Architecture of a Subscriber Access Dual-Stack Network 389
Terms Used in IPv6 Subscriber Management Documentation 390
IPv6 Addressing Requirements for a Subscriber Access Network 392
Alternatives to Using a Global IPv6 Address on the CPE WAN Link 393

Chapter 43 Introduction to IPv6 Addresses 395

IPv6 Addressing Overview 395
IPv6 Notation 395
IPv6 Prefixes 396
IPv6 Address Types 396
Unicast Addresses 397
Multicast Addresses 397
Anycast Addresses 397

Chapter 44 Using NDRA to Provide IPv6 WAN Link Addressing 399

Using NDRA to Provide IPv6 WAN Link Addressing Overview 399
IPv6 Neighbor Discovery Protocol Overview 399
Neighbor Discovery Messages 400
How NDRA Works in a Subscriber Access Network 400
Methods for Obtaining IPv6 Prefixes for NDRA 401
Using AAA RADIUS Server to Obtain IPv6 Prefixes for NDRA 401
Duplicate Prefix Protection for NDRA 402

Chapter 45 Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing 403

Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview 403
Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA 403
Using a AAA RADIUS Server to Obtain IPv6 Addresses for DHCPv6 IA_NA 403

Chapter 46 Using DHCPv6 Prefix Delegation to Provide Subscriber LAN Addressing 405

Using DHCPv6 Prefix Delegation Overview 405
Using a Delegated Prefix on the CPE Loopback Interface 407
DHCPv6 Prefix Delegation over PPPoE 407
Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation 408
Using a AAA RADIUS Server to Obtain IPv6 Prefixes for Prefix Delegation 408

Chapter 47 Using Both DHCPv6 IA_NA and DHCPv6 Prefix Delegation to Provide IPv6 WAN Link and Subscriber LAN Addressing 409

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview 409
Lease Times and Session Timeouts for DHCPv6 IA_NA and DHCPv6 Prefix Delegation 409
DHCPv6 Options in a DHCPv6 Multiple Address Environment 410

Copyright © 2017, Juniper Networks, Inc.

Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA . . . 410
Address Pools for DHCPv6 Prefix Delegation and DHCPv6 IA_NA . . . 411
Using a AAA RADIUS Server to Obtain IPv6 Addresses and Prefixes . . . 411
Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment . . 411
Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface . . . 412

Chapter 48 Designs for IPv6 Addressing in a Subscriber Access Network . . . 413

Selecting the Type of Addressing Used on the CPE . . . 413
Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link . . . 413
Selecting the Method of Assigning Global IPv6 Addresses to Subscribers . . . 414
Selecting the Method of Obtaining IPv6 Prefixes . . . 414
Using a AAA RADIUS Server to Obtain Global IPv6 Addresses and IPv6 Prefixes . . . 415
Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation . . . 415
Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation . . . 416
Design 3: IPv6 Addressing with NDRA . . . 417
Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix . . . 418

Chapter 49 Dual-Stack Access Models in a DHCP Network . . . 419

IPv4 and IPv6 Dual Stack in a DHCP Access Network . . . 419
Support for Demultiplexing Interfaces . . . 419
AAA Service Framework in a Dual Stack over a DHCP Access Network . . . 420
Collection of Accounting Statistics in a DHCP Access Network . . . 420
Change of Authorization (CoA) . . . 421
Dual-Stack Interface Stack in a DHCP Wholesale Network . . . 421
Single-Session DHCP Dual-Stack Overview . . . 422
Configuring Single-Session DHCP Dual-Stack Support . . . 423
Verifying and Managing DHCP Dual-Stack Configuration . . . 425

Chapter 50 Dual-Stack Access Models in a PPPoE Network . . . 427

IPv4 and IPv6 Dual Stack in a PPPoE Access Network . . . 427
Support for Demultiplexing Interfaces . . . 428
Determining the Status of CPE in a PPPoE Access Network . . . 428
IPV6 Address Provisioning in the PPPoE Access Network . . . 428
Authentication in a PPPoE Access Network . . . 429
Negotiation of Network Control Protocols When Authorized Addresses Are Unavailable . . . 429
Shared IPv4 and IPv6 Service Sessions on PPP Access Networks . . . 430
Accounting for Shared IPv4 and IPv6 Service Sessions . . . 430
Deactivating Shared IPv4 and IPv6 Service Sessions . . . 430
AAA Service Framework in a Dual Stack over a PPPoE Access Network . . . 430
Collection of Accounting Statistics in a PPPoE Access Network . . . 431
Change of Authorization (CoA) . . . 432

Copyright © 2017, Juniper Networks, Inc.

RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers 432
Accounting Messages for PPPoE Using NDRA Prefixes 433
Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes 438
Suppressing Accounting Information That Comes from AAA 444
Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address 444

Chapter 51 **Configuring DHCPv6 Local Server 447**

DHCPv6 Local Server Overview 448
Enabling DHCPv6 Rapid Commit Support 449
Preventing Binding of Clients That Do Not Support Reconfigure Messages 450
Configuring the DUID Type Supported by DHCPv6 Servers 451
Example: Extended DHCPv6 Local Server Configuration 452

Chapter 52 **Configuring DHCPv6 Relay Agent 455**

DHCPv6 Relay Agent Overview 455
DHCPv6 Relay Agent Options 456
Configuring DHCPv6 Relay Agent Options 456
Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets 458
Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets 459

Chapter 53 **Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation 461**

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation 461
On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview 462
IPv4 Address Negotiation for Static PPP Subscribers 462
IPv4 Address Release for Static PPP Subscribers 464
On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview 465
IPv4 Address Negotiation for Dynamic PPP Subscribers 465
IPv4 Address Release for Dynamic PPP Subscribers 467
IPCP Negotiation with Optional Peer IP Address 468
How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled 468
Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers 469
Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers 470
Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers 470
Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes 471
Enabling Unisphere IPv4 Release Control VSA in RADIUS Messages 471

Chapter 54 **Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network 473**

Best Practice: Static PPPoE Interfaces with NDRA 473
Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network 474

Copyright © 2017, Juniper Networks, Inc.

Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA . . . 474
Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 . . . 475
Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles . . . 476
Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network . . . 477

Chapter 55 Configuring Dual Stack for PPPoE Access Networks in Which DHCP Is Used . . . 479

Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE . . . 479
Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network . . . 480

Chapter 56 Configuring Dual Stack for PPPoE Access Networks in Which NDRA Is Used . . . 483

Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network . . . 483
Configuring a Static PPPoE Logical Interface for NDRA . . . 485
Configuring an Address-Assignment Pool Used for Router Advertisements . . . 486
Configuring Duplicate IPv6 Prefix Protection for Router Advertisement . . . 487

Chapter 57 Configuring Address Assignment Pools for DHCPv6 . . . 489

Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation . . . 489
Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA . . . 490
Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation . . . 490

Chapter 58 Configuring the Dynamic Router Advertisement Protocol . . . 493

Dynamic Router Advertisement Configuration Overview . . . 493

Chapter 59 Examples: IPv4 and IPv6 Dual-Stack Designs . . . 495

Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE . . . 495
Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE . . . 518
Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE . . . 542

Chapter 60 Monitoring and Managing Dual Stack Subscribers . . . 561

Monitoring Active Subscriber Sessions . . . 561
Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance . . . 562
Monitoring Dynamic Subscriber Sessions . . . 562
Monitoring Address Pools Used for Subscribers . . . 563
Monitoring Specific Subscriber Sessions . . . 564
Monitoring the Status of the PPPoE Logical Interface . . . 565
Monitoring Service Sessions for Subscribers . . . 566
Monitoring PPP Options Negotiated with the Remote Peer . . . 566
Monitoring the RADIUS Attribute Used for NDRA . . . 567

Copyright © 2017, Juniper Networks, Inc.

Chapter 61 Monitoring and Managing DHCPv6 569
Monitoring Address Bindings on the DHCPv6 Local Server 569
Verifying and Managing DHCPv6 Local Server Configuration 570
Verifying and Managing DHCPv6 Relay Configuration 570

Part 4 Configuring Junos OS Enhanced Subscriber Management

Chapter 62 Configuring Junos OS Enhanced Subscriber Management 573
Junos OS Enhanced Subscriber Management Overview 573
Configuring Junos OS Enhanced Subscriber Management 575
Verifying and Managing Junos OS Enhanced Subscriber Management 577
Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview 578
Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview 579

Part 5 Configuring Address-Assignment Pools for Subscriber Management

Chapter 63 Configuring Address-Assignment Pools for Dynamic and Static Addresses 583
Address-Assignment Pools Overview 583
Address-Assignment Pools Licensing Requirements 584
Configuring Address-Assignment Pools 585
Example: Configuring an Address-Assignment Pool 586
Configuring an Address-Assignment Pool Name and Addresses 587
Configuring a Named Address Range for Dynamic Address Assignment 588
Configuring Address-Assignment Pool Usage Threshold Traps 588
Configuring Address-Assignment Pool Linking 589
Configuring Address-Assignment Pool Hold-Down 590
Configuring DHCP Local Address Pool Rapid Drain 591
Configuring Static Address Assignment 594
Configuring Duplicate IPv4 Address Protection for AAA 595

Part 6 Configuring DNS Addresses for Subscriber Management

Chapter 64 Configuring DNS Address Assignments and Session Options 599
DNS Name Server Address Overview 599
Configuring DNS Name Server Addresses for Subscriber Management 600
Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses 602
Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment 603
DNS Resolver for IPv6 DNS Overview 604
Configuring a DNS Server Address for IPv6 Hosts 604

Copyright © 2017, Juniper Networks, Inc.

Part 7 Configuring CLI-Based Subscriber Services

Chapter 65 Configuring CLI-Activated Subscriber Services 609

CLI-Activated Subscriber Services 609
Activating and Deactivating Subscriber Services Locally with the CLI 610
Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers 613
Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers 614
Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber 614

Chapter 66 Configuring Subscriber Services with Multiple Instances 617

Subscriber Services with Multiple Instances Overview 617
Subscriber Service Instances and Service Parameters 617
CLI Deactivation of Subscriber Services with Multiple Instances 618
Subscriber Services with Multiple Instances in RADIUS Accounting Messages 618
Deactivating a Single Instance of a Subscriber Service with Multiple Instances 619
Deactivating All Instances of a Subscriber Service with Multiple Instances 622

Chapter 67 Monitoring and Managing Subscriber Services 625

Verifying and Managing Subscriber Services with Multiple Instances 625

Part 8 Configuring ANCP and the ANCP Agent for Subscribers

Chapter 68 Configuring ANCP Agent Neighbors and Operations 629

ANCP and the ANCP Agent Overview 629
Overview 630
Topology Discovery 630
Subscriber Services 631
ANCP Interfaces and Access Loop Circuit Identifiers 632
Mapping Access Lines to Interfaces and Interface Sets 632
ANCP Neighbors 634
Partitions 635
Adjacency Update Messages 636
Generic Response Messages and Result Codes 636
ANCP Operations in Different Network Configurations 638
1:1 and N:1 Traffic Shaping Models 639
Business Services Traffic Shaping Model 641
ANCP Network Using N:1 and 1:1 Configuration Models without Interface Sets 642
Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets 643
ANCP Network Using N:1 Configuration Model with Interface Sets 644
Sequence of ANCP Events: Static VLAN Interfaces over Ethernet with Interface Sets 645
ANCP Network Using 1:1 Configuration Model with Interface Sets 646

Copyright © 2017, Juniper Networks, Inc.

Sequence of ANCP Events: Static VLAN Demux Interfaces over Aggregated Ethernet with Interface Sets . . . 647
Configuring the ANCP Agent . . . 648
Configuring ANCP Neighbors . . . 649
Associating an Access Node with Subscribers for ANCP Agent Operations . . . 651
Specifying the Interval Between ANCP Adjacency Messages . . . 652
Specifying the Maximum Number of Discovery Table Entries . . . 652
Configuring the ANCP Agent for Backward Compatibility . . . 653
Specifying How Long Processes Wait for the ANCP Agent Restart to Complete . . . 654
Configuring the ANCP Agent to Learn ANCP Partition IDs . . . 654
Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet . . . 655

Chapter 69 **Configuring the ANCP Agent Traffic and CoS . . . 677**

Traffic Rate Reporting and Adjustment by the ANCP Agent . . . 677
Overview . . . 677
Traffic Rate Adjustment . . . 679
Recommended Traffic Shaping Rates . . . 679
ANCP Agent Keepalives for CoS . . . 680
Preservation of CoS Shaping Across ANCP Agent Restarts . . . 680
Configuring the ANCP Agent to Report Traffic Rates to CoS . . . 681
Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces . . . 683
Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates . . . 685

Chapter 70 **Configuring the ANCP Agent and AAA . . . 687**

ANCP Agent Interactions with AAA . . . 687
ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes . . . 689
Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages . . . 691
Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications . . . 692

Chapter 71 **Monitoring and Managing ANCP for Subscriber Access . . . 695**

Triggering ANCP OAM to Test the Local Loop . . . 695
Verifying and Monitoring ANCP Neighbors . . . 696
Clearing ANCP Neighbors . . . 697
Verifying and Monitoring ANCP Subscribers . . . 697
Clearing ANCP Subscribers . . . 698
Verifying and Monitoring CoS for ANCP Subscribers . . . 698
Clearing and Verifying ANCP Statistics . . . 699

Copyright © 2017, Juniper Networks, Inc.

Part 9 Configuring the Diameter Base Protocol and its Applications

Chapter 72 Configuring the Diameter Base Protocol 703

Diameter Base Protocol Overview 703
Messages Used by Diameter Applications 706
Diameter AVPs and Diameter Applications 711
Configuring Diameter 723
Configuring the Origin Attributes of the Diameter Instance 724
Configuring Diameter Peers 724
Configuring the Diameter Transport 725
Configuring Diameter Network Elements 726

Chapter 73 Configuring Gx-Plus for Provisioning Subscribers 729

Gx-Plus for Provisioning Subscribers Overview 729
Understanding Gx-Plus Interactions Between the Router and the PCRF 731
Subscriber Login 731
Fault Tolerance and Event Notification 733
PCRF-Generated Discovery 735
Subscriber Accounting 735
Subscriber Audit 735
Subscriber Logout 736
Configuring Gx-Plus 736
Configuring the Gx-Plus Partition 737
Configuring Gx-Plus Global Attributes 738
Provisioning Subscribers with Gx-Plus 739

Chapter 74 Configuring 3GPP Policy and Charging Control 741

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting 741
Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers 744
Wireline Access Environment 744
Junos OS Environment 746
Understanding Gx Interactions Between the Router and the PCRF 748
Subscriber Login 748
Subscriber Update 752
Subscriber Logout 753
Subscriber Disconnect 754
Connectivity Fault Recovery 755
Understanding Gy Interactions Between the Router and the OCS 756
First Interrogation to the OCS 756
Intermediate Interrogation to the OCS 758
Final Interrogation to the OCS 759
Connectivity Fault Recovery 760
Abort Session Requests 761
Understanding Rule Interactions Between the PCRF, PCEF, and OCS 761
Login Interactions 762
Update Interactions 762
Quota Expiration and Validity-Time Interactions 763
Connection and Monitoring Interactions 763

Copyright © 2017, Juniper Networks, Inc.

Logout Interactions . . . 765
Understanding Upstream and Downstream Messages for the PCRF . . . 766
Common Upstream Messages . . . 767
Event-Timestamp AVP . . . 767
Charging Rules Installation Notifications . . . 767
Event Trigger Commands . . . 768
Common Downstream Messages . . . 768
Charging Rule Installation Commands . . . 768
Charging Rule Removal Commands . . . 768
Event Trigger Commands . . . 769
Configuring the OCS Partition . . . 769
Configuring the PCRF Partition . . . 773
Configuring OCS Global Parameters . . . 776

Chapter 75 Configuring NASREQ for Authentication and Authorization . . . 779

Diameter Network Access Server Application (NASREQ) . . . 779
Configuring the Diameter Network Access Server Application (NASREQ) . . . 781

Chapter 76 Configuring JSRC in Subscriber Access Networks . . . 783

Juniper Networks Session and Resource Control (SRC) and JSRC Overview . . . 783
Hardware Requirements for JSRC for Subscriber Access . . . 784
Understanding JSRC-SAE Interactions . . . 785
Subscriber Login . . . 785
Subscriber Service Activation and Deactivation . . . 786
Subscriber Resynchronization . . . 786
Subscriber Session Terminated by the SAE . . . 786
Statistics Collection and Reporting per Service Rule . . . 787
Subscriber Logout . . . 787
Configuring JSRC . . . 787
Configuring the JSRC Partition . . . 788
Assigning a Partition to JSRC . . . 789
Authorizing Subscribers with JSRC . . . 789
Provisioning Subscribers with JSRC . . . 790

Chapter 77 Excluding Diameter AVPs from JSRC Messages . . . 791

Excluding AVPs from Diameter Messages for JSRC . . . 791

Chapter 78 Configuring Service Accounting with JSRC . . . 793

Service Accounting with JSRC . . . 793
Configuring Service Accounting with JSRC . . . 794

Chapter 79 Configuring Subscribers on Static Interfaces . . . 797

Subscribers on Static Interfaces Overview . . . 797
Configuring Subscribers over Static Interfaces . . . 800
Example: Configuring Static Subscribers for Subscriber Access . . . 801

Chapter 80 Configuring the Static Subscribers Global Profile . . . 805

Specifying the Static Subscriber Global Access Profile . . . 805
Specifying the Static Subscriber Global Dynamic Profile . . . 805
Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers . . . 806

Copyright © 2017, Juniper Networks, Inc.

Configuring the Static Subscriber Global Authentication Password 807
Configuring the Static Subscriber Global Username 807

Chapter 81 **Configuring the Static Subscribers Group Profile 809**

Creating a Static Subscriber Group 809
Specifying the Static Subscriber Group Access Profile 810
Specifying the Static Subscriber Group Dynamic Profile 810
Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group 811
Configuring the Static Subscriber Group Authentication Password 811
Configuring the Static Subscriber Group Username 812

Chapter 82 **Monitoring and Managing Diameter Information for Subscriber Access . . 815**

Verifying Diameter Node, Instance, and Route Information 815
Verifying and Managing Diameter Function Information 816
Verifying and Managing Diameter Peer Information 817
Verifying Diameter Network Element Information 818

Chapter 83 **Monitoring and Managing Subscriber Information on Static Interfaces . . 819**

Forcing a Static Subscriber to Be Logged Out 819
Resetting the State of an Interface for Static Subscriber Login 819
Forcing a Group of Static Subscribers to Be Logged Out 820
Resetting the State of an Interface Group for Static Subscriber Login 820
Verifying Information about Subscriber Sessions on Static Interfaces 820

Part 10 Troubleshooting

Chapter 84 **Configuring AAA Testing and Troubleshooting 823**

AAA Configuration Testing and Troubleshooting 823
Testing a Subscriber AAA Configuration 823

Chapter 85 **Tracing Extended DHCP Operations 829**

Tracing Extended DHCP Operations 829
Configuring the Extended DHCP Log Filename 831
Configuring the Number and Size of Extended DHCP Log Files 831
Configuring Access to the Extended DHCP Log File 832
Configuring a Regular Expression for Extended DHCP Messages to Be Logged 832
Configuring the Extended DHCP Tracing Flags 832
Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged 833
Tracing Extended DHCP Operations for Specific Interfaces 834

Chapter 86 **Configuring Subscriber Management Database Log Files 837**

Configuring the Number and Size of Subscriber Management Database Log Files 837
Configuring Access to the Subscriber Management Database Log File 838
Configuring a Regular Expression for Subscriber Management Database Messages to Be Logged 838

Copyright © 2017, Juniper Networks, Inc.

Chapter 87 Configuring Subscriber Management Database Trace Flags and Operations 839
Tracing Subscriber Management Database Operations for Subscriber Access 839
Configuring the Subscriber Management Database Trace Log Filename 840
Configuring the Subscriber Management Database Tracing Flags 840
Chapter 88 Configuring Subscriber Management Session Database Log Files 843
Configuring the Number and Size of Subscriber Management Session Database Replication Log Files 843
Configuring Access to the Subscriber Management Session Database Replication Log File 844
Configuring a Regular Expression for Subscriber Management Session Database Replication Messages to Be Logged 844
Chapter 89 Configuring Subscriber Management Session Database Trace Flags and Operations 845
Tracing Subscriber Management Session Database Replication Operations for Subscriber Access 845
Configuring the Subscriber Management Session Database Replication Trace Log Filename 846
Configuring the Subscriber Management Session Database Replication Tracing Flags 846
Chapter 90 Configuring Diameter Base Protocol Log Files 849
Configuring the Number and Size of Diameter Base Protocol Log Files 849
Configuring Access to the Diameter Base Protocol Log File 850
Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged 850
Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged 850
Chapter 91 Configuring Diameter Base Protocol Trace Flags and Operations 853
Tracing Diameter Base Protocol Processes for Subscriber Access 853
Configuring the Diameter Base Protocol Trace Log Filename 854
Configuring the Diameter Base Protocol Tracing Flags 854
Chapter 92 Troubleshooting Diameter Networks 857
Troubleshooting Diameter Network Configuration 857
Troubleshooting Diameter Network Connectivity 858
Chapter 93 Configuring ANCP Log Files 859
Configuring the Number and Size of ANCP Log Files 859
Configuring Access to the ANCP Log File 860
Configuring a Regular Expression for ANCP Messages to Be Logged 860
Configuring the Severity Level to Filter Which ANCP Messages Are Logged . . . 860
Chapter 94 Configuring ANCP Trace Flags and Operations 863
Tracing ANCP Agent Operations for Subscriber Access 863
Configuring the ANCP Trace Log Filename 864
Configuring the ANCP Tracing Flags 864

Copyright © 2017, Juniper Networks, Inc.

Chapter 95 Configuring General Authentication Service Log Files 867
Configuring the Number and Size of General Authentication Service Processes Log Files 867
Configuring Access to the Log File 868
Configuring a Regular Expression for Lines to Be Logged 868
Chapter 96 Configuring General Authentication Service Trace Flags and Operations 869
Tracing General Authentication Service Processes 869
Configuring the General Authentication Service Processes Trace Log Filename 870
Configuring the Number and Size of General Authentication Service Processes Log Files 870
Configuring Access to the Log File 871
Configuring a Regular Expression for Lines to Be Logged 871
Configuring Subscriber Filtering for General Authentication Service Trace Operations 871
Configuring the Trace Operation 872
Configuring the General Authentication Service Processes Trace Log Filename 873
Configuring the Trace Operation 874
Configuring Subscriber Filtering for General Authentication Service Trace Operations 874
Chapter 97 Configuring Static Subscriber Interfaces Log Files 877
Configuring the Number and Size of Static Subscribers Log Files 877
Configuring Access to the Static Subscribers Log File 878
Configuring a Regular Expression for Static Subscriber Messages to Be Logged 878
Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged 878
Chapter 98 Configuring Static Subscriber Interfaces Trace Flags and Operations . . . 881
Tracing Static Subscriber Operations 881
Configuring the Static Subscribers Trace Log Filename 882
Configuring the Static Subscribers Tracing Flags 882
Chapter 99 Overriding PCRF Session Control to Troubleshoot a Session or Services 885
Disabling PCRF Control of a Subscriber Session 885
Chapter 100 Contacting Juniper Technical Support 887
Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support 887
Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support 889

Copyright © 2017, Juniper Networks, Inc.

Part 11 Configuration Statements and Operational Commands

Chapter 101 Configuration Statements 895

aaa-logical-system (Domain Map) 909
aaa-routing-instance (Domain Map) 910
abated-utilization (Address-Assignment Pools) 911
abated-utilization-v6 (Address-Assignment Pools) 912
accept-max-tcp-connections (System Process) 913
access-identifier 914
access-loop-id-local 914
access-profile 915
access-profile (Extensible Subscriber Services Manager) 916
access-profile (Domain Map) 916
access-profile (Static Subscribers) 917
access-profile-name (Duplicate Accounting) 918
accounting (Access Profile) 919
accounting (Service Accounting) 920
accounting-backup-options (Access Profile) 920
accounting-order (Service Accounting) 921
accounting-port 922
accounting-retry (RADIUS) 923
accounting-server 924
accounting-session-id-format 924
accounting-stop-on-access-deny 925
accounting-stop-on-failure 926
accounting-timeout (RADIUS) 927
active-drain (Address Assignment Pools) 928
active-server-group 929
actual-transit-statistics (Dynamic Profiles) 930
address (Diameter Peer) 930
address (Diameter Transport) 931
address-assignment (Address-Assignment Pools) 932
address-change-immediate-update 933
address-pool (Domain Map) 934
address-protection 935
adjacency-timer 936
adsl-bytes 937
adsl2-bytes 938
adsl2-plus-bytes 939
advisory-options (Traffic Shaping) 940
aggregate-clients (DHCP Relay Agent) 941
aggregate-clients (Static Subscribers) 943
allow-bulk-leasequery (DHCP Local Server) 944
allow-leasequery (DHCP Local Server) 945
always-write-giaddr 946
always-write-option-82 947
ancp 948
ancp-speed-change-immediate-update (ANCP) 949
asymmetric-lease-time (DHCP Overrides) 950

Copyright © 2017, Juniper Networks, Inc.

asymmetric-prefix-lease-time (DHCP Overrides) . . . 951
attempts (DHCP Local Server) . . . 952
attempts (DHCP Relay Agent) . . . 953
attributes . . . 954
attributes (JSRC Attributes) . . . 955
authentication (DHCP Local Server) . . . 956
authentication (DHCP Relay Agent) . . . 957
authentication (Static Subscribers) . . . 958
authentication-order . . . 959
authentication-server . . . 960
authorization-order . . . 961
autonomous (Dynamic Router Advertisement) . . . 962
boot-file . . . 962
boot-server . . . 963
bulk-leasequery (DHCP Relay Agent) . . . 964
called-station-id (OCS Partition) . . . 965
calling-station-id-delimiter (Subscriber Management) . . . 966
calling-station-id-format (Subscriber Management) . . . 967
chap-challenge-in-request-authenticator (Subscriber Management) . . . 968
charging-id (OCS Partition) . . . 969
charging-service-list . . . 970
circuit-id (DHCP Relay Agent) . . . 971
circuit-id (Address-Assignment Pools) . . . 973
circuit-type (DHCP Local Server) . . . 974
circuit-type (DHCP Relay Agent) . . . 975
clear-on-abort (DHCP Local Server) . . . 976
client-accounting-algorithm . . . 977
client-authentication-algorithm . . . 978
client-discover-match (DHCP Local Server) . . . 980
client-discover-match (DHCP Relay Agent) . . . 981
client-id (DHCP Local Server) . . . 982
client-id (DHCP Relay Agent) . . . 983
client-idle-timeout . . . 984
client-idle-timeout-ingress-only . . . 985
client-negotiation-match (DHCPv6 Local Server) . . . 986
client-negotiation-match (DHCPv6 Relay Agent) . . . 987
client-session-timeout . . . 988
commit-interval . . . 989
coa-dynamic-variable-validation . . . 990
coa-immediate-update . . . 990
coa-no-override service-class-attribute . . . 991
configuration-database (Enhanced Subscriber Management) . . . 991
concurrent-data-sessions . . . 992
connect-actively . . . 992
current-hop-limit (Dynamic Router Advertisement) . . . 993
database-replication (Subscriber Session Database) . . . 993
default-action (DHCP Relay Agent Option) . . . 994
default-lifetime (Dynamic Router Advertisement) . . . 995
delay-advertise (DHCPv6) . . . 996

Copyright © 2017, Juniper Networks, Inc.

delay-authentication (DHCP Relay Agent) . . . 998
delay-offer (DHCPv4) . . . 999
delegated-pool (DHCP Local Server) . . . 1000
delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) . . . 1001
delimiter (DHCP Local Server) . . . 1002
delimiter (DHCP Relay Agent) . . . 1004
delimiter (Domain Map) . . . 1005
destination (Diameter Network Element) . . . 1006
destination-host . . . 1006
destination-host (Gx-Plus) . . . 1007
destination-host (OCS Partition) . . . 1007
destination-host (PCRF Partition) . . . 1008
destination-realm (JSRC) . . . 1008
destination-realm (Gx-Plus) . . . 1009
destination-realm (OCS Partition) . . . 1010
destination-realm (PCRF Partition) . . . 1011
dhcp-attributes (Address-Assignment Pools) . . . 1012
dhcp-local-server . . . 1014
dhcp-relay . . . 1020
dhcp-service . . . 1029
dhcpv6 (DHCP Local Server) . . . 1030
dhcpv6 (DHCP Relay Agent) . . . 1034
diameter . . . 1040
diameter-instance (JSRC) . . . 1041
diameter-instance (Diameter Applications) . . . 1042
dictionary . . . 1043
disable . . . 1043
disable (Extensible Subscriber Services Manager) . . . 1044
disable-relay . . . 1045
dne-origin (Diameter Network Element) . . . 1046
dns-server . . . 1047
dns-server-address (Dynamic Profiles) . . . 1048
domain (Domain Map) . . . 1049
domain-name (DHCP Local Server) . . . 1050
domain-name (DHCP Relay Agent) . . . 1052
domain-name (Address-Assignment Pools) . . . 1053
domain-name (Static Subscribers) . . . 1054
domain-name-server (Routing Instances and Access Profiles) . . . 1055
domain-name-server-inet (Routing Instances and Access Profiles) . . . 1056
domain-name-server-inet6 (Routing Instances and Access Profiles) . . . 1057
downstream-rate (Traffic Shaping) . . . 1058
draining (Diameter Applications) . . . 1059
draining-response-timeout (Diameter Applications) . . . 1060
drop (DHCP Relay Agent Option) . . . 1061
dual-stack (DHCP Relay Agent Overrides) . . . 1062
dual-stack-group (DHCP Relay Agent) . . . 1063
duplication (Access Profile) . . . 1064
duplication-filter (Access Profile) . . . 1065
duplication-vrf (Duplicate Accounting) . . . 1066

Copyright © 2017, Juniper Networks, Inc.

dynamic-profile (DHCP Local Server) . . . 1067
dynamic-profile (DHCP Relay Agent) . . . 1068
dynamic-profile (Domain Map) . . . 1069
dynamic-profile (Static Subscribers) . . . 1070
enable . . . 1071
enable (Enhanced Subscriber Management) . . . 1071
equals (DHCP Relay Agent) . . . 1072
ethernet-port-type-virtual . . . 1073
exceed-action . . . 1073
exclude (JSRC Attributes) . . . 1074
exclude (RADIUS) . . . 1075
external-authority . . . 1080
failover (System Process) . . . 1081
family (Address-Assignment Pools) . . . 1082
fast-lookup-filter . . . 1083
final-response-timeout (OCS Partition) . . . 1084
force-continue (OCS Partition) . . . 1085
forward-only (DHCP Relay Agent Option) . . . 1086
forward-only (DHCP Relay Agent) . . . 1087
forward-only-replies (DHCP Relay Agent) . . . 1088
forwarding (Diameter Network Element) . . . 1089
function (Diameter Network Element) . . . 1090
function (Diameter Route) . . . 1091
ggsn-address (OCS Partition) . . . 1092
ggsn-mcc-mnc (OCS Partition) . . . 1093
global (Gx-Plus) . . . 1094
global (OCS) . . . 1095
global (PCRF) . . . 1096
grace-period . . . 1097
group (DHCP Local Server) . . . 1098
group (DHCP Relay Agent) . . . 1101
group (Static Subscribers) . . . 1105
gsmp-syn-timeout (ANCP) . . . 1106
gsmp-syn-wait (ANCP) . . . 1107
gx-plus (Gx-Plus) . . . 1108
hardware-address . . . 1109
high-utilization (Address-Assignment Pools) . . . 1109
high-utilization-v6 (Address-Assignment Pools) . . . 1110
hold-down (Address-Assignment Pools) . . . 1111
host (Address-Assignment Pools) . . . 1112
host-name (dhcp-relay) . . . 1113
host-name (dhcpv6) . . . 1114
ietf-mode . . . 1115
ignore . . . 1116
immediate-update . . . 1117
include-ipv6 (Gx-Plus) . . . 1117
include-irb-and-l2 . . . 1118
include-l2-interface-for-irb (DHCP Relay Agent) . . . 1120
include-option-82 (DHCP Local Server) . . . 1121

Copyright © 2017, Juniper Networks, Inc.

interface (DHCP Local Server) . . . 1122
interface (DHCP Relay Agent) . . . 1124
interface (Dynamic Router Advertisement) . . . 1126
interface (Static Subscriber Group) . . . 1127
interface (Static Subscriber Username) . . . 1128
interface-client-limit (DHCP Local Server) . . . 1129
interface-client-limit (DHCP Relay Agent) . . . 1131
interface-delete (Subscriber Management or DHCP Client Management) . . . 1132
interface-description (DHCP Local Server) . . . 1133
interface-description (DHCP Relay Agent) . . . 1134
interface-description-format . . . 1135
interface-name (DHCP Local Server) . . . 1136
interface-name (DHCP Relay Agent) . . . 1137
interface-mib (Enhanced Subscriber Management) . . . 1138
interface-set (ANCP) . . . 1139
interface-traceoptions (DHCP) . . . 1140
interfaces (ANCP) . . . 1142
ip-address . . . 1143
ip-address-change-notify . . . 1144
ip-address-first . . . 1145
ip-can-type (PCRF Partition) . . . 1146
ipcp-suggest-dns-option . . . 1147
jsrc (JSRC) . . . 1148
jsrc (Access Profile) . . . 1149
jsrc-partition . . . 1150
juniper-dsl-attributes . . . 1151
layer2-unicast-replies . . . 1152
keep-incoming-circuit-id (DHCP Relay Agent) . . . 1153
keep-incoming-interface-id (DHCP Relay Agent) . . . 1154
keep-incoming-remote-id (DHCP Relay Agent) . . . 1155
leasequery (DHCP Relay Agent) . . . 1156
lease-time-threshold (DHCP Local Server and DHCP Relay Agent) . . . 1157
lease-time-validation (DHCP Local Server and DHCP Relay Agent) . . . 1158
limit . . . 1158
link (Address-Assignment Pools) . . . 1159
local (Flat-File Access Profile) . . . 1160
local-decision (PCRF Partition) . . . 1161
local-server-group (DHCP Relay Agent Option) . . . 1162
location (dhcp-relay) . . . 1163
location (dhcpv6) . . . 1164
logical-interface-unit-range . . . 1165
logical-system (Diameter Peer) . . . 1166
logical-system (Diameter Transport) . . . 1167
logical-system (Subscriber Secure Policy) . . . 1168
logical-system-name (Static Subscribers) . . . 1169
logical-system-name (DHCP Local Server) . . . 1170
logical-system-name (DHCP Relay Agent) . . . 1171
logout-response-timeout (PCRF Partition) . . . 1172
ltv-syslog-interval (System Process) . . . 1173

Copyright © 2017, Juniper Networks, Inc.

mac-address (DHCP Local Server) . . . 1174
mac-address (DHCP Relay Agent) . . . 1175
maintain-subscriber (Subscriber Management) . . . 1176
managed-configuration (Dynamic Router Advertisement) . . . 1177
map (Domain Map) . . . 1178
mask (Domain Map) . . . 1179
max-advertisement-interval (Dynamic Router Advertisement) . . . 1180
max-connections (DHCP Local Server) . . . 1181
max-data-sessions-per-subscriber . . . 1182
max-db-size (Enhanced Subscriber Management) . . . 1183
max-empty-replies (DHCP Local Server) . . . 1184
max-outstanding-requests . . . 1185
max-outstanding-requests (Diameter Applications) . . . 1186
max-pending-accounting-stops (Access Profile) . . . 1187
max-withhold-time (Access Profile) . . . 1187
maximum-discovery-table-entries . . . 1188
maximum-helper-restart-time . . . 1189
maximum-lease-time . . . 1190
maximum-subscribers . . . 1191
metric (Diameter Route) . . . 1191
metric (Domain Map) . . . 1192
min-advertisement-interval (Dynamic Router Advertisement) . . . 1192
multi-address-embedded-option-response (DHCP Local Server) . . . 1193
name-server . . . 1194
nas-identifier . . . 1194
nas-port-extended-format . . . 1195
nas-port-extended-format (Interfaces) . . . 1197
nas-port-id-delimiter (Subscriber Management) . . . 1198
nas-port-id-format (Subscriber Management) . . . 1199
nas-port-options (RADIUS Options) . . . 1201
nas-port-type (Subscriber Management) . . . 1203
nas-port-type (RADIUS Options) . . . 1205
nasreq (Diameter Application) . . . 1207
neighbor (Define ANCP) . . . 1208
neighbor-discovery-router-advertisement (Address-Assignment Pools) . . . 1209
netbios-node-type . . . 1209
network . . . 1210
network-element (Diameter Base Protocol) . . . 1211
network-services . . . 1212
no-bind-on-request (DHCP Relay Agent) . . . 1214
no-unsolicited-ra (Enhanced Subscriber Management) . . . 1215
no-vlan-interface-name . . . 1216
ocs (Diameter Applications) . . . 1218
on-demand-ip-address . . . 1220
on-link (Dynamic Router Advertisement) . . . 1221
option . . . 1222
option-60 (DHCP Local Server) . . . 1223
option-60 (DHCP Relay Agent) . . . 1224
option-82 (DHCP Relay Agent) . . . 1225

Copyright © 2017, Juniper Networks, Inc.

option-82 (DHCP Local Server Authentication) . . . 1226
option-82 (DHCP Local Server Pool Matching) . . . 1227
option-82 (Address-Assignment Pools) . . . 1228
option-match . . . 1229
option-number (DHCP Relay Agent Option) . . . 1230
options (Access Profile) . . . 1231
order . . . 1233
origin (Diameter Base Protocol) . . . 1234
other-bytes . . . 1235
other-overhead-adjust . . . 1236
other-stateful-configuration (Dynamic Router Advertisement) . . . 1237
overhead-accounting (ANCP) . . . 1238
overrides (DHCP Local Server) . . . 1239
overrides (DHCP Relay Agent) . . . 1241
overrides (Enhanced Subscriber Management) . . . 1243
override-password (Domain Map) . . . 1244
padn (Domain Map) . . . 1244
parse-direction (Domain Map) . . . 1245
parse-order (Domain Map) . . . 1246
partition . . . 1247
partition (Gx-Plus) . . . 1248
partition (NASREQ Diameter Application) . . . 1249
partition (OCS) . . . 1250
partition (PCRF) . . . 1252
password (Static Subscribers) . . . 1254
password (DHCP Local Server) . . . 1255
password (DHCP Relay Agent) . . . 1257
pcrf (Diameter Applications) . . . 1259
peer (Diameter Base Protocol) . . . 1260
peer (Diameter Network Element) . . . 1261
peer-ip-address-optional . . . 1262
peer-origin (Diameter Peer) . . . 1263
pool (Address-Assignment Pools) . . . 1264
pool (DHCP Local Server Overrides) . . . 1265
pool-match-order . . . 1266
port . . . 1267
port (Diameter Peer) . . . 1267
pre-ietf-mode . . . 1268
preauthentication-order (Access Profile) . . . 1269
preauthentication-port . . . 1270
preauthentication-secret . . . 1271
preauthentication-server (Access Profile) . . . 1272
preferred-lifetime (Address-Assignment Pools) . . . 1273
preferred-lifetime (Dynamic Router Advertisement) . . . 1274
prefix (DHCP Relay Agent) . . . 1275
prefix (Address-Assignment Pools) . . . 1276
prefix (Dynamic Router Advertisement) . . . 1277
priority (Diameter Peer) . . . 1277
profile (Access) . . . 1278

Copyright © 2017, Juniper Networks, Inc.

process-inform . . . 1283
protocols (Dynamic Profiles) . . . 1285
provisioning-order Diameter Applications) . . . 1287
proxy-mode . . . 1288
qos-adjust . . . 1289
qos-adjust-adsl . . . 1290
qos-adjust-adsl2 . . . 1291
qos-adjust-adsl2-plus . . . 1292
qos-adjust-other . . . 1293
qos-adjust-sdsl . . . 1294
qos-adjust-vdsl . . . 1295
qos-adjust-vdsl2 . . . 1296
radius (Access Profile) . . . 1297
radius-disconnect (DHCP Local Server) . . . 1299
radius-flow-tap . . . 1301
radius-options (Edit Access) . . . 1302
radius-options (Interfaces) . . . 1303
radius-server . . . 1304
range (Address-Assignment Pools) . . . 1305
rapid-commit (DHCPv6 Local Server) . . . 1306
reachable-time (Dynamic Router Advertisement) . . . 1307
realm-delimiter (Domain Map) . . . 1307
realm-parse-direction (Domain Map) . . . 1308
reconfigure (DHCP Local Server) . . . 1309
relay-agent-interface-id (DHCP Local Server) . . . 1310
relay-agent-interface-id (DHCPv6 Relay Agent) . . . 1311
relay-agent-interface-id (DHCPv6 Relay Agent Username) . . . 1312
relay-agent-remote-id (DHCP Local Server) . . . 1313
relay-agent-remote-id (DHCPv6 Relay Agent) . . . 1314
relay-agent-remote-id (DHCPv6 Relay Agent Username) . . . 1315
relay-agent-subscriber-id (DHCP Local Server) . . . 1316
relay-agent-subscriber-id (DHCPv6 Relay Agent) . . . 1317
relay-option (DHCP Relay Agent) . . . 1318
relay-option-vendor-specific (dhcpv6) . . . 1319
relay-option-82 . . . 1320
relay-server-group (DHCP Relay Agent Option) . . . 1321
relay-source . . . 1322
remote-id . . . 1323
remote-id (DHCP Relay Agent) . . . 1324
remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) . . . 1326
replace-ip-source-with . . . 1327
report-interface-descriptions (Edit Access) . . . 1328
report-local-rule (PCRF Partition) . . . 1329
report-resource-allocation (PCRF Partition) . . . 1330
report-successful-resource-allocation (PCRF Partition) . . . 1332
request-max-tcp-connections (System Process) . . . 1333
request network-access aaa replay pending-accounting-stops . . . 1334
request network-access aaa subscriber add session-id . . . 1335
request network-access aaa subscriber delete session-id . . . 1337

Copyright © 2017, Juniper Networks, Inc.

request network-access aaa subscriber modify session-id 1339
request-rate 1340
requested-ip-network-match (DHCP Local Server) 1341
restricted-requestor (DHCP Local Server) 1342
retransmit-timer (Dynamic Router Advertisement) 1343
retry 1344
revert-interval 1345
route (Diameter Network Element) 1346
router (Address-Assignment Pools) 1346
router-advertisement (Dynamic Profiles) 1347
routing-instance 1347
routing-instance (Subscriber Secure Policy) 1348
routing-instance (Diameter Peer) 1348
routing-instance (Diameter Transport) 1349
routing-instance-name (DHCP Local Server) 1350
routing-instance-name (DHCP Relay Agent) 1352
routing-instance-name (Static Subscribers) 1353
routing-services (Enhanced Subscriber Management) 1354
sdsl-bytes 1355
sdsl-overhead-adjust 1356
secret 1357
send-acct-status-on-config-change (Access Profile) 1358
send-dyn-subscription-indicator (PCRF Partition) 1359
send-network-family-indicator (PCRF Partition) 1361
send-origin-state-id (Diameter Applications) 1362
send-release-on-delete (DHCP Relay Agent) 1363
server-duid-type (DHCP Local Server) 1364
server-group 1365
server-identifier (Address-Assignment Pools) 1366
server-id-override 1367
server-response-time (DHCP Relay Agent) 1368
service (Service Accounting) 1369
service-activation (Access Profile) 1370
service-context-id (OCS) 1372
service-profile (DHCP Local Server) 1373
service-profile (DHCP Relay Agent) 1374
session-options 1375
shmlog (Shared Memory Log) 1376
sip-server-address 1378
sip-server-domain-name 1378
smg-service (Enhanced Subscriber Management) 1379
source-address 1380
stacked-vlan-ranges (RADIUS Options) 1381
starts-with (DHCP Relay Agent Option) 1382
static-subscribers 1383
statistics (Access Profile) 1384
statistics (Service Accounting) 1385
strict (DHCP Local Server) 1386
strip-domain (Domain Map) 1387

Copyright © 2017, Juniper Networks, Inc.

strip-username (Domain Map) . . . 1387
subscriber-packet-idle-timeout . . . 1388
subscriber-management (Subscriber Management) . . . 1389
subscriber-profile . . . 1390
subscription-id-data-include (PCRF Partition) . . . 1391
subscription-id-type (PCRF Partition) . . . 1393
t1-percentage (Address-Assignment Pools) . . . 1394
t1-renewal-time (Address-Assignment Pools) . . . 1395
t2-percentage (Address-Assignment Pools) . . . 1396
t2-rebinding-time (Address-Assignment Pools) . . . 1397
target-logical-system (Domain Map) . . . 1398
target-routing-instance (Domain Map) . . . 1399
terminate-code . . . 1400
tftp-server . . . 1401
timeout (RADIUS) . . . 1402
timeout (DHCP Local Server) . . . 1403
timeout (DHCP Local Server Leasequery) . . . 1404
timeout (DHCP Relay Agent Leasequery) . . . 1405
timeout-grace (RADIUS Options) . . . 1406
token (DHCP Local Server) . . . 1407
trace (DHCP Local Server) . . . 1408
trace (DHCP Relay Agent) . . . 1409
traceoptions (ANCP) . . . 1410
traceoptions (DHCP) . . . 1412
traceoptions (Diameter Base Protocol) . . . 1415
traceoptions (Extensible Subscriber Services Manager) . . . 1417
traceoptions (Enhanced Subscriber Management) . . . 1418
traceoptions (General Authentication Service) . . . 1421
traceoptions (Static Subscribers) . . . 1423
traceoptions (Subscriber Management) . . . 1425
traceoptions (Subscriber Session Database Replication) . . . 1427
transport (Diameter Base Protocol) . . . 1428
transport (Diameter Peer) . . . 1429
traps . . . 1430
trigger (DHCP Local Server) . . . 1431
trigger (DHCP Relay Agent Leasequery) . . . 1432
trio-flow-offload . . . 1433
trust-option-82 . . . 1434
tunnel-profile (Domain Map) . . . 1435
underlying-interface (ANCP) . . . 1435
unique-nas-port . . . 1436
update-interval . . . 1437
update-interval (Service Accounting) . . . 1438
update-response-timeout (PCRF Partition) . . . 1439
upstream-rate (Traffic Shaping) . . . 1440
use-interface-description . . . 1441
user-name-include (OCS Partition) . . . 1443
use-option-82 . . . 1445
use-primary (DHCP Relay Agent) . . . 1446

Copyright © 2017, Juniper Networks, Inc.

use-vlan-id . . . 1448
user-prefix (DHCP Local Server) . . . 1450
user-prefix (DHCP Relay Agent) . . . 1452
user-prefix (Static Subscribers) . . . 1453
username-include (DHCP Local Server) . . . 1454
username-include (DHCP Relay Agent) . . . 1456
username-include (Static Subscribers) . . . 1458
valid-lifetime (Dynamic Router Advertisement) . . . 1459
valid-lifetime (Address-Assignment Pools) . . . 1460
vdsl-bytes . . . 1461
vdsl-overhead-adjust . . . 1462
vdsl2-bytes . . . 1463
vdsl2-overhead-adjust . . . 1464
vendor-specific (dhcp-relay) . . . 1465
violation-action (DHCP Local Server and DHCP Relay Agent) . . . 1466
vlan-nas-port-stacked-format . . . 1467
vlan-ranges (RADIUS Options) . . . 1468
vrf-name (Duplicate Accounting) . . . 1469
wait-for-acct-on-ack (Access Profile) . . . 1469
wins-server (Access) . . . 1470

Chapter 102 Operational Commands . . . 1471

clear ancp neighbor . . . 1475
clear ancp statistics . . . 1477
clear ancp subscriber . . . 1479
clear dhcp relay binding . . . 1481
clear dhcp relay statistics . . . 1484
clear dhcp server binding . . . 1487
clear dhcp server statistics . . . 1490
clear dhcpv6 relay binding . . . 1492
clear dhcpv6 relay statistics . . . 1495
clear dhcpv6 server binding . . . 1497
clear dhcpv6 server statistics . . . 1499
clear diameter function statistics . . . 1500
clear diameter peer . . . 1501
clear extensible-subscriber-services counters . . . 1502
clear extensible-subscriber-services sessions . . . 1503
clear ipv6 router-advertisement . . . 1504
clear network-access aaa statistics . . . 1505
clear network-access aaa subscriber . . . 1507
clear network-access gx-plus replay . . . 1509
clear network-access gx-plus statistics . . . 1510
clear network-access ocs statistics . . . 1511
clear network-access pcrf . . . 1512
request services subscribers clear . . . 1513
clear system subscriber-management statistics . . . 1514
request ancp oam interface . . . 1515
request ancp oam neighbor . . . 1516
request ancp oam port-down . . . 1518

Copyright © 2017, Juniper Networks, Inc.

request ancp oam port-up . . . 1520
request dhcp relay bulk-leasequery . . . 1522
request dhcp relay leasequery . . . 1524
request dhcp server reconfigure . . . 1526
request dhcpv6 server reconfigure . . . 1528
request dhcpv6 relay bulk-leasequery . . . 1530
request dhcpv6 relay leasequery . . . 1532
request network-access aaa subscriber set session-id . . . 1534
request services extensible-subscriber-services reload-dictionary . . . 1536
request services static-subscribers login group . . . 1537
request services static-subscribers logout group . . . 1538
request services static-subscribers login interface . . . 1539
request services static-subscribers logout interface . . . 1540
request system reboot . . . 1541
restart extensible-subscriber-services . . . 1546
set request services subscribers . . . 1547
show accounting pending-accounting-stops . . . 1548
show ancp cos . . . 1552
show ancp neighbor . . . 1557
show ancp statistics . . . 1565
show ancp subscriber . . . 1570
show ancp summary . . . 1576
show ancp summary neighbor . . . 1578
show ancp summary subscriber . . . 1580
show class-of-service interface . . . 1582
show class-of-service interface-set . . . 1611
show class-of-service scheduler-map . . . 1614
show class-of-service traffic-control-profile . . . 1617
show database-replication statistics . . . 1621
show database-replication summary . . . 1623
show dhcp relay binding . . . 1625
show dhcp relay statistics . . . 1631
show dhcp server binding . . . 1635
show dhcp server statistics . . . 1641
show dhcpv6 relay binding . . . 1645
show dhcpv6 relay statistics . . . 1653
show dhcpv6 server binding . . . 1657
show dhcpv6 server statistics . . . 1663
show diameter . . . 1666
show diameter function . . . 1672
show diameter function statistics . . . 1676
show diameter instance . . . 1679
show diameter network-element . . . 1681
show diameter network-element map . . . 1684
show diameter peer . . . 1687
show diameter peer map . . . 1692
show diameter peer statistics . . . 1695
show diameter route . . . 1699
show dynamic-profile session . . . 1701

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services accounting 1706
show extensible-subscriber-services counters 1709
show extensible-subscriber-services debug-information 1711
show extensible-subscriber-services dictionary 1712
show extensible-subscriber-services dictionary attributes 1716
show extensible-subscriber-services dictionary services 1719
show extensible-subscriber-services sessions 1722
show extensible-subscriber-services service 1724
show ipv6 router-advertisement 1725
show network-access aaa accounting 1728
show network-access aaa radius-servers 1729
show network-access aaa statistics 1736
show network-access aaa statistics authentication 1745
show network-access aaa statistics pending-accounting-stops 1748
show network-access aaa statistics preauthentication 1749
show network-access aaa subscribers 1751
show network-access aaa subscribers session-id 1754
show network-access aaa terminate-code 1761
show network-access address-assignment pool 1766
show network-access domain-map 1768
show network-access gx-plus 1770
show network-access nasreq statistics 1772
show network-access ocs 1776
show network-access pcrf 1778
show ppp address-pool 1781
show route 1783
show route extensive 1800
show static-subscribers sessions 1818
show subscribers 1820
show subscribers summary 1846
show system subscriber-management route 1852
show system subscriber-management statistics 1857
show system subscriber-management summary 1863
test aaa authd-lite user 1866
test aaa dhcp user 1869
test aaa ppp user 1873

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

List of Figures

Part 1 Configuring AAA for Subscriber Management

Chapter 5 **Configuring RADIUS Accounting for Subscriber Access 97**

Figure 1: Topology with Loss of Access to Accounting Server 102

Chapter 10 **Receiving DHCP Options From a RADIUS Server 141**

Figure 2: DHCP Options Data Flow 144

Part 3 Configuring IPv6 for Subscriber Management

Chapter 42 **Migrating to IPv6 Using IPv4 and IPv6 Dual Stack 389**

Figure 3: IPv4 and IPv6 Dual-Stack Architecture in a Subscriber Access Network 390

Figure 4: IPv6 Address Requirements in a Subscriber Access Network 392

Chapter 46 **Using DHCPv6 Prefix Delegation to Provide Subscriber LAN Addressing 405**

Figure 5: Delegated Addressing in a Dual-Stack Network Using DHCPv6 406

Chapter 48 **Designs for IPv6 Addressing in a Subscriber Access Network 413**

Figure 6: Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation 416

Figure 7: Subscriber Access Network with NDRA and DHCPv6 Prefix Delegation 417

Figure 8: Subscriber Access Network with NDRA 417

Figure 9: Subscriber Access Network with DHCPv6 Prefix Delegation 418

Chapter 49 **Dual-Stack Access Models in a DHCP Network 419**

Figure 10: Dual-Stack Interface Stack over a DHCP Access Network 419

Figure 11: AAA Service Framework in a Dual Stack over a DHCP Access Network 420

Figure 12: Dual-Stack Interface Stack in a DHCP Wholesale Network 421

Chapter 50 **Dual-Stack Access Models in a PPPoE Network 427**

Figure 13: Dual-Stack Interface Stack over a PPPoE Access Network 428

Figure 14: Dual-Stack Aggregated Ethernet Stack over a PPPoE Access Network 428

Figure 15: AAA Service Framework in a Dual Stack over a PPPoE Access Network 431

Chapter 59 **Examples: IPv4 and IPv6 Dual-Stack Designs 495**

Figure 16: PPPoE Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation 496

Copyright © 2017, Juniper Networks, Inc.

Figure 17: PPPoE Subscriber Access Network with ND/RA and DHCPv6 Prefix Delegation 519
Figure 18: PPPoE Subscriber Access Network with NDRA 542

Part 8 Configuring ANCP and the ANCP Agent for Subscribers

Chapter 68 Configuring ANCP Agent Neighbors and Operations 629

Figure 19: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model) . . 642
Figure 20: Sample ANCP Topology with Interface Sets (N:1 Model) 644
Figure 21: Sample ANCP Topology with Interface Sets (1:1 Model) 646
Figure 22: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet 656

Part 9 Configuring the Diameter Base Protocol and its Applications

Chapter 74 Configuring 3GPP Policy and Charging Control 741

Figure 23: 3GPP PCC Architecture Overview 742
Figure 24: 3GPP PCC Architecture Overview 744
Figure 25: 3GPP PCC Architecture Overview 761

Copyright © 2017, Juniper Networks, Inc.

List of Tables

About the Documentation xlv

Table 1: Notice Icons xlvii
Table 2: Text and Syntax Conventions xlviii

Part 1 Configuring AAA for Subscriber Management

Chapter 2 Configuring RADIUS Attributes and Juniper Networks VSAs 13

Table 3: Supported RADIUS IETF Attributes 14
Table 4: Supported Juniper Networks VSAs 22
Table 5: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs 33
Table 6: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs 39
Table 7: Attributes That Can Be Ignored in RADIUS Access-Accept Messages . . . 44
Table 8: Attributes That Can Be Excluded from RADIUS Messages 44
Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables 50
Table 10: DSL Forum VSAs 57
Table 11: DSL Forum VSAs—Supported RADIUS Messages 58
Table 12: Microsoft Vendor-Specific RADIUS Attributes for DNS Server Addresses 60

Chapter 3 Configuring RADIUS NAS-Port Attributes and Options 61

Table 13: RADIUS NAS-Port-Type Values 76

Chapter 5 Configuring RADIUS Accounting for Subscriber Access 97

Table 14: RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting 98
Table 15: Duplicate RADIUS Accounting Reporting 99
Table 16: Juniper Networks VSAs Used for Per-Service Session Accounting 109

Chapter 10 Receiving DHCP Options From a RADIUS Server 141

Table 17: Unsupported Opaque DHCP Options 146
Table 18: DHCP Options Description 147

Chapter 11 Configuring RADIUS Logical Line Identification 149

Table 19: RADIUS Attributes for LLID Preauthentication Requests 151

Chapter 12 Provisioning Subscriber Services Using Cisco VSAs 155

Table 20: Effective Service Accounting Method with CLI and RADIUS Configurations 156
Table 21: Effective Service Interim Accounting Interval with CLI and RADIUS Configurations 156

Copyright © 2017, Juniper Networks, Inc.

Chapter 13 Configuring Domain Maps for Subscriber Management 159
Table 22: Domain Map Options and Parameters 161
Table 23: Precedence Rules for Applying Access Profiles 167
Table 24: Precedence Rules for Determining the Address Pool to Use 168
Table 25: Precedence Rules for Applying Dynamic Profiles 169
Chapter 14 Configuring Dynamic Service Activation for Subscriber Access 177
Table 26: Identification Attributes 180
Table 27: Session Attributes 180
Table 28: Juniper Network VSAs Used for Service Thresholds 182
Table 29: Error-Cause Codes (RADIUS Attribute 101) 184
Chapter 15 Configuring Termination Causes for Protocols 187
Table 30: RFC-Defined Code Values and Termination Causes 188
Table 31: Default Mapping Between AAA Termination Causes and Code Values 191
Table 32: Default Mapping Between DHCP Termination Causes and Code Values 193
Table 33: Default Mapping Between L2TP Termination Causes and Code Values 194
Table 34: Default Mapping Between PPP Termination Causes and Code Values 214
Table 35: Default Mapping Between VLAN Termination Causes and Code Values 222
Part 2 Configuring DHCP for Subscriber Management
Chapter 18 Using DHCP Overview 233
Table 36: Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server 235
Table 37: Information in Authentication Grant 237
Chapter 20 Assigning IP Addresses 249
Table 38: DHCP Attributes 249
Table 39: DHCPv6 Attributes 250
Chapter 21 Configuring Lease Times for IP Addresses 257
Table 40: Lease-Time Violation Event Logging 259
Chapter 28 Forcing Dynamic Reconfiguration of Clients From a DHCP Local Server 305
Table 41: Action Taken for Events That Occur During a Reconfiguration 307
Chapter 29 Conserving IP Addresses Using DHCP Auto Logout 313
Table 42: DHCP Relay Agent Option 82 Value for Auto Logout 316
Part 3 Configuring IPv6 for Subscriber Management
Chapter 42 Migrating to IPv6 Using IPv4 and IPv6 Dual Stack 389
Table 43: IPv6 Subscriber Management Terms 390
Chapter 48 Designs for IPv6 Addressing in a Subscriber Access Network 413

Copyright © 2017, Juniper Networks, Inc.

Table 44: Choosing the Global IPv6 Address Provisioning Method for the WAN Link . . . 414
Table 45: RADIUS Attributes Used to Obtain Global IPv6 Addresses and IPv6 Prefixes . . . 415

Chapter 51 **Configuring DHCPv6 Local Server . . . 447**

Table 46: RADIUS Attributes and VSAs for DHCPv6 Local Server . . . 448

Chapter 59 **Examples: IPv4 and IPv6 Dual-Stack Designs . . . 495**

Table 47: Configuration Components Used in Dual Stack with DHCPv6 IA_NA and DHCPv6 Prefix Delegation . . . 496
Table 48: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation . . . 520
Table 49: Configuration Components Used in Dual Stack with ND/RA and DHCPv6 Prefix Delegation . . . 543

Part 7 Configuring CLI-Based Subscriber Services

Chapter 66 **Configuring Subscriber Services with Multiple Instances . . . 617**

Table 50: Subscriber Services and Service Parameters in RADIUS Accounting Messages . . . 619

Part 8 Configuring ANCP and the ANCP Agent for Subscribers

Chapter 68 **Configuring ANCP Agent Neighbors and Operations . . . 629**

Table 51: ANCP Failure Result Codes . . . 637
Table 52: ACI Mapping by Interface Type for the ANCP 1:1 Model . . . 639
Table 53: ACI Mapping by Interface Set Type for the ANCP N:1 Access Model . . 640
Table 54: Configuration Components used in ANCP N:1 Topology Example with Interface Sets . . . 657

Chapter 70 **Configuring the ANCP Agent and AAA . . . 687**

Table 55: Mapping ANCP DSL Attributes to Juniper Networks DSL VSAs and DSL Forum VSAs . . . 690

Part 9 Configuring the Diameter Base Protocol and its Applications

Chapter 72 **Configuring the Diameter Base Protocol . . . 703**

Table 56: Diameter Messages and Diameter Applications . . . 707
Table 57: Standard Diameter AVPs . . . 711
Table 58: Juniper Networks Diameter AVPs . . . 718
Table 59: Tekelec Diameter AVPs . . . 722

Chapter 73 **Configuring Gx-Plus for Provisioning Subscribers . . . 729**

Table 60: Differences Between Gx-Plus and Junos OS Terminology . . . 729
Table 61: Router Events, Router Actions, and PCRF Actions . . . 734

Chapter 74 **Configuring 3GPP Policy and Charging Control . . . 741**

Table 62: Functionality Comparison Between PCRF and PCEF . . . 743
Table 63: Allocated Juniper AVPs . . . 745
Table 64: Family Indicator AVP . . . 746
Table 65: Initial Set of Fixed Juniper-Parameter VSAs . . . 747

Copyright © 2017, Juniper Networks, Inc.

Table 66: Result-Code-AVP Categories 749
Table 67: CCA-GX-I Response Processing 750
Table 68: Result-Code-AVP Categories 753
Table 69: Result-Code-AVP Categories 754
Table 70: Result-Code-AVP Categories 755
Table 71: Event-Timestamp AVP Message Definition 767

Part 11 Configuration Statements and Operational Commands

Chapter 101 Configuration Statements 895

Table 72: Service Activation/Deactivation Error Messages 1335
Table 73: Service Activation/Deactivation Error Messages 1337
Table 74: Service Activation/Deactivation Error Messages 1339
Table 75: Allocated Juniper AVPs 1359
Table 76: Network Family Indicator AVP 1361

Chapter 102 Operational Commands 1471

Table 77: clear dhcp relay statistics Output Fields 1485
Table 78: Service Activation/Deactivation Error Messages 1534
Table 79: show accounting pending-accounting-stops Output Fields 1548
Table 80: show ancp cos Output Fields 1552
Table 81: show ancp neighbor Output Fields 1557
Table 82: show ancp statistics Output Fields 1565
Table 83: show ancp subscriber Output Fields 1571
Table 84: show ancp summary Output Fields 1576
Table 85: show ancp summary neighbor Output Fields 1578
Table 86: show ancp summary subscriber Output Fields 1580
Table 87: show class-of-service interface Output Fields 1583
Table 88: show class-of-service interface-set Output Fields 1611
Table 89: show class-of-service scheduler-map Output Fields 1614
Table 90: show class-of-service traffic-control-profile Output Fields 1617
Table 91: show database-replication statistics Output Fields 1621
Table 92: show database-replication summary Output Fields 1623
Table 93: show dhcp relay binding Output Fields 1626
Table 94: show dhcp relay statistics Output Fields 1632
Table 95: show dhcp server binding Output Fields 1636
Table 96: show dhcp server statistics Output Fields 1642
Table 97: show dhcpv6 relay binding Output Fields 1646
Table 98: show dhcpv6 relay statistics Output Fields 1654
Table 99: show dhcpv6 server binding Output Fields 1658
Table 100: show dhcpv6 server statistics Output Fields 1664
Table 101: show diameter Output Fields 1666
Table 102: show diameter function Output Fields 1672
Table 103: show diameter function statistics Output Fields 1676
Table 104: show diameter instance Output Fields 1679
Table 105: show diameter network-element Output Fields 1681
Table 106: show diameter network-element map Output Fields 1684
Table 107: show diameter peer Output Fields 1687
Table 108: show diameter peer map Output Fields 1692
Table 109: show diameter peer statistics Output Fields 1695

Copyright © 2017, Juniper Networks, Inc.

Table 110: show diameter route Output Fields 1699
Table 111: show extensible-subscriber-services accounting Output Fields 1706
Table 112: show extensible-subscriber-services counters Output Fields 1709
Table 113: show extensible-subscriber-services debug-information Output Fields 1711
Table 114: show extensible-subscriber-services dictionary Output Fields 1712
Table 115: show extensible-subscriber-services dictionary attributes Output Fields 1716
Table 116: show extensible-subscriber-services dictionary services Output Fields 1719
Table 117: show extensible-subscriber-services sessions Output Fields 1722
Table 118: show extensible-subscriber-services service Output Fields 1724
Table 119: show ipv6 router-advertisement Output Fields 1726
Table 120: show network-access aaa accounting Output Fields 1728
Table 121: show network-access aaa radius-servers Output Fields 1729
Table 122: show network-access aaa statistics Output Fields 1737
Table 123: show network-access aaa statistics authentication Output Fields . . 1745
Table 124: show network-access aaa statistics pending-accounting-stops Output Fields 1748
Table 125: show network-access aaa statistics preauthentication Output Fields 1749
Table 126: show network-access aaa subscribers Output Fields 1751
Table 127: show network-access aaa subscribers session-id Output Fields . . . 1755
Table 128: show network-access aaa terminate-code Output Fields 1762
Table 129: show network-access address-assignment pool Output Fields 1766
Table 130: show network-access domain-map Output Fields 1768
Table 131: show network-access gx-plus Output Fields 1770
Table 132: show network-access nasreq statistics Output Fields 1772
Table 133: show network-access ocs Output Fields 1776
Table 134: show network-access pcrf Output Fields 1779
Table 135: show ppp address-pool Output Fields 1781
Table 136: show route Output Fields 1784
Table 137: show route extensive Output Fields 1800
Table 138: show static-subscribers sessions Output Fields 1818
Table 139: show subscribers Output Fields 1824
Table 140: show subscribers summary Output Fields 1847
Table 141: show system subscriber-management route Output Fields 1853
Table 142: show system subscriber-management statistics Output Fields 1857
Table 143: show system subscriber-management summary Output Fields . . . 1863

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

About the Documentation

- Documentation and Release Notes on page xlv
- Supported Platforms on page xlv
- Using the Examples in This Manual on page xlv
- Documentation Conventions on page xlvii
- Documentation Feedback on page xlix
- Requesting Technical Support on page xlix

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

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

- MX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

Copyright © 2017, Juniper Networks, Inc.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```


Copyright © 2017, Juniper Networks, Inc.

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 on page xlvii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xlviii defines the text and syntax conventions used in this guide.

Copyright © 2017, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(***string1* **|** *string2* **|** *string3***)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** *community-ids* **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Copyright © 2017, Juniper Networks, Inc.

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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

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

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

Copyright © 2017, Juniper Networks, Inc.

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

Copyright © 2017, Juniper Networks, Inc.

PART 1

Configuring AAA for Subscriber Management

- AAA and RADIUS for Subscriber Access Overview on page 3
- Configuring RADIUS Attributes and Juniper Networks VSAs on page 13
- Configuring RADIUS NAS-Port Attributes and Options on page 61
- Configuring RADIUS Authentication for Subscriber Access on page 91
- Configuring RADIUS Accounting for Subscriber Access on page 97
- Configuring Routers and RADIUS Servers for Subscriber Access on page 113
- Configuring Access Profiles for Subscriber Access on page 123
- Configuring the Effect of Service Activation Failures on Subscriber Login on page 125
- Configuring Session Options for Subscriber Access on page 133
- Receiving DHCP Options From a RADIUS Server on page 141
- Configuring RADIUS Logical Line Identification on page 149
- Provisioning Subscriber Services Using Cisco VSAs on page 155
- Configuring Domain Maps for Subscriber Management on page 159
- Configuring Dynamic Service Activation for Subscriber Access on page 177
- Configuring Termination Causes for Protocols on page 187
- Configuring Extensible Subscriber Services Manager on page 225
- Monitoring and Managing AAA Information for Subscriber Access on page 227

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 1

AAA and RADIUS for Subscriber Access Overview

- AAA Service Framework Overview on page 3
- RADIUS Server Options for Subscriber Access on page 4
- Global RADIUS Options for Subscriber Access on page 8
- Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview on page 8

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.

Copyright © 2017, Juniper Networks, Inc.

- 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 Router or Switch Interaction with RADIUS Servers on page 113
- RADIUS Acct-On and Acct-Off Messages on page 105
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Address-Assignment Pools Overview on page 583
- RADIUS Accounting Statistics for Subscriber Access Overview on page 97
- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- *Subscriber Secure Policy Overview*

RADIUS Server Options for Subscriber Access

You can specify options that the router uses when communicating with RADIUS authentication and accounting servers for subscriber access.

Copyright © 2017, Juniper Networks, Inc.

The following list describes the RADIUS options you can configure:

- **access-loop-id-local**—The Agent-Remote-Id and Agent-Circuit-Id generated locally when these values are not present in the client database. The interface description of the logical interface is used as the Agent-Remote-Id and the interface description portion of the NAS-Port-Id using the format **<underlying-interface-name>:<outer-tag>-<inner-tag>** is used as the Agent-Circuit-Id.

  NOTE: The NAS-Port-Id format changes (established by **[set access profile *profile-name* radius options interface-description-format]**) are applied before generating the Agent-Circuit-Id.

  The NAS-Port-Id format (established by **[set access profile *profile-name* radius options interface-description-format]**) leverages the locally generated Agent-Remote-Id and Agent-Circuit-Id.

- **accounting-session-id-format**—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**
- **calling-station-id-delimiter**—The character that the router uses as the separator between concatenated values in the Calling-Station-Id string (RADIUS attribute 31).
- **calling-station-id-format**—Optional information that the router includes in the Calling-Station-Id (RADIUS attribute 31).
- **client-accounting-algorithm** and **client-authentication-algorithm**—The method the router uses to access RADIUS accounting and RADIUS authentication servers. You can specify the following methods:
  - **direct**—The default method, in which there is no load balancing. For example, in the direct method, the router always accesses **server1** (the primary server) first, and uses **server2** and **server3** as backup servers.
  - **round-robin**—The method that provides load balancing by rotating router requests among the list of configured RADIUS servers. For example, if three RADIUS servers are configured to support the router, the router sends the first request to **server1**, and uses **server2** and **server3** as backup servers. The router then sends the second request to **server2**, and uses **server3** and **server1** as backups.

    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.

Copyright © 2017, Juniper Networks, Inc.

- **coa-dynamic-variable-validation**—The optional method that the router uses when processing CoA requests that include changes to a client profile dynamic variable that cannot be applied. The optional configuration specifies that when a CoA operation is unable to apply a requested change to a client profile dynamic variable, subscriber management does not apply any changes to client profile dynamic variables in the CoA request and then responds with a NACK. In the default method, 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.
- **ethernet-port-type-virtual**—The physical port type of **virtual** that the router uses 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.
- **interface-description-format**—The information that is excluded from the interface description that the router passes to RADIUS for inclusion in the RADIUS attribute 87 (NAS-Port-Id). By default, the router includes both the **subinterface** and the **adapter** in the interface description. You can specify:
  - **exclude-adapter**—Exclude the adapter.
  - **exclude-subinterface**—Exclude the subinterface.
- **nas-identifier**—The value for the client RADIUS attribute 32 (NAS-Identifier), which is used for authentication and accounting requests. You can specify a string in the range 1 through 64 characters.
- **nas-port-extended-format**—The extended format for RADIUS attribute 5 (NAS-Port) and for the width of the fields in the NAS-Port attribute that the RADIUS client uses. You can specify:
  - **adapter-width *width***—Number of bits in the adapter field.
  - **port-width *width***—Number of bits in the port field.
  - **pw-width**—Number of bits in the pseudowire field.
  - **slot-width *width***—Number of bits in the slot field.
  - **stacked-vlan-width *width***—Number of bits in the SVLAN ID field.
  - **vlan-width *width***—Number of bits in the VLAN ID field.

  NOTE: **The total of the widths must not exceed 32 bits, or the configuration fails.**

  You can configure an extended format for the NAS-Port attribute for both Ethernet subscribers and ATM subscribers. For ATM subscribers, you can specify:
  - **adapter-width**—Number of bits in the ATM adapter field, in the range 0 through 32
  - **port-width**—Number of bits in the ATM port field, in the range 0 through 32
  - **slot-width**—Number of bits in the ATM slot field, in the range 0 through 32

Copyright © 2017, Juniper Networks, Inc.

  - **vci-width**—Number of bits in the ATM virtual circuit identifier (VCI) field, in the range 0 through 32
  - **vpi-width**—Number of bits in the ATM virtual path identifier (VPI) field, in the range 0 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.

- **nas-port-id-delimiter**—The character used as the separator between values in the NAS-Port-Id string.
- **nas-port-id-format**—Optional information included in RADIUS attribute 87 (NAS-Port-Id).
- **nas-port-type**—The port type used to authenticate subscribers.
- **revert-interval**—The number of seconds that the router waits after a server has become unreachable. The router rechecks the connection to the server when the **revert-interval** expires. If the server is then reachable, it is used in accordance with the order of the server list. You can configure from 0 (off) through 604800 seconds. The default is 60 seconds.
- **service-activation**—Setting that determines 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 (**dynamic-profile**) or in Extensible Subscriber Services Manager (ESSM) operation scripts (**extensible-service**):
  - **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. This is the default behavior for the **extensible-service** service type.
  - **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. This is the default value for the **dynamic-profile** service type.
- **vlan-nas-port-stacked-format**—The format that turns off RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

Global RADIUS Options for Subscriber Access

You can specify options that the router uses when communicating with all configured RADIUS servers for subscriber access.

The following list describes the global RADIUS options you can configure:

- **revert-interval**—The number of seconds that the router waits after a server has become unreachable. The router rechecks the connection to the server when the **revert-interval** expires. If the server is then reachable, it is used in accordance with the order of the server list. You can configure from 0 (off) through 604800 seconds. The default is 60 seconds.
- **request-rate**—The number of requests per second that the router can send 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. You can configure from 100 through 4000 requests per second. The default is 500 requests per second.
- **timeout-grace**—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. You can configure a duration from 0 through 30 seconds; the default period is 10 seconds.

Related Documentation

- Configuring RADIUS Options for Subscriber Access Globally on page 120

Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview

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.

You can enable or disable storage and reporting of interface descriptions as follows:

- To enable storing and reporting of interface descriptions, include the **report-interface-descriptions** statement at the **[edit access]** hierarchy level.
- To disable storing and reporting of interface descriptions, include the **radius attributes exclude** statement at the **[edit access profile *profile-name*** hierarchy level.

Copyright © 2017, Juniper Networks, Inc.

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.

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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
                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;
            }
        }
```


Copyright © 2017, Juniper Networks, Inc.

```
        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;
            }
        }
    }
}
```

Interface Descriptions on Aggregated Ethernet Physical Interfaces

For aggregated Ethernet interfaces, the interface description on the aggregated Ethernet interface, for example AE0 or AE1, serves as the physical interface description.

Interface Descriptions on a Combination of Dynamic and Static Interfaces

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 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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
                dynamic-profile vlan-prof {
                    accept inet;
                    ranges {
                        any;
                    }
                }
                authentication {
                    password $ABC123;
                    username-include {
                        user-prefix rich;
                    }
                }
            }
        }
    }
```

Related Documentation

- report-interface-descriptions (Edit Access) on page 1328
- exclude on page 1075
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 2

Configuring RADIUS Attributes and Juniper Networks VSAs

- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 13
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 14
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 33
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 49
- DSL Forum Vendor-Specific Attributes on page 56
- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS on page 58
- RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses on page 59

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

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 14
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- DSL Forum Vendor-Specific Attributes on page 56
- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 33
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 49

RADIUS IETF Attributes Supported by the AAA Service Framework

Table 3 on page 14 describes the RADIUS IETF attributes that the Junos OS AAA Service Framework supports.

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 3: 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
4	NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.	No

Copyright © 2017, Juniper Networks, Inc.

Table 3: Supported RADIUS IETF Attributes *(continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
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
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
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
18	Reply-Message	• Text that may be displayed to the user. • Only the first instance of this attribute is used.	No

Copyright © 2017, Juniper Networks, Inc.

Table 3: 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>]**	Yes
25	Class	Arbitrary value that the NAS includes in all accounting packets for the user if supplied by the RADIUS server.	No
27	Session-Timeout	Maximum number of consecutive seconds of service to be provided to the user before termination of the session.	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
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

Copyright © 2017, Juniper Networks, Inc.

Table 3: Supported RADIUS IETF Attributes *(continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
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
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 3: Supported RADIUS IETF Attributes *(continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 3: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 3: 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, Ip: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
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
123	Delegated-IPv6-Prefix	IPv6 prefix that is delegated to the user.	No

Copyright © 2017, Juniper Networks, Inc.

Table 3: Supported RADIUS IETF Attributes *(continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
242	Ascend-Data-Filter	Binary data that specifies RADIUS policy definitions.	Yes

Related Documentation

- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 33
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39

Juniper Networks VSAs Supported by the AAA Service Framework

Table 4 on page 22 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).

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.

Copyright © 2017, Juniper Networks, Inc.

Table 4: Supported Juniper Networks VSAs

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-1	Virtual-Router	Client logical system:routing instance name. Allowed only from AAA server for default logical system:routing instance. When this VSA is not included in the subscriber profile, the routing instance assigned to the subscriber—the one in which the subscriber session comes up—varies by subscriber type. For DHCP and PPPoE subscribers, it is the default routing instance. For L2TP tunnel subscribers, it is the routing instance in which the tunnel resides, whether default or non-default. If the tunnel routing instance is not default and you want the L2TP session to be in the default routing instance, you must use the Virtual-Router VSA to set the desired routing instance.	string: *logical system:routing instance*	No
26-4	Primary-DNS	Client DNS address negotiated during IPCP.	integer: 4-byte *primary-dns-address*	No
26-5	Secondary-DNS	Client DNS address negotiated during IPCP	integer: 4-byte *secondary-dns-address*	No
26-6	Primary-WINS	Client WINS (NBNS) address negotiated during IPCP.	integer: 4-byte *primary-wins-address*	No
26-7	Secondary-WINS	Client WINS (NBNS) address negotiated during IPCP.	integer: 4-byte *secondary-wins-address*	No
26-8	Tunnel-Virtual-Router	Virtual router name for tunnel connection.	string: *tunnel-virtual-router*	No
26-9	Tunnel-Password	Tunnel password in cleartext. Do not use both this VSA and the standard RADIUS attribute Tunnel-Password [69]. We recommend that you use the standard attribute because the password is encrypted when that attribute is used.	string: *tunnel-password*	No
26-10	Ingress-Policy-Name	Input policy name to apply to client interface.	string: *input-policy-name*	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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
26-23	IGMP-Enable	Whether IGMP is enabled or disabled on a client interface.	integer: • 0=disable • 1=enable	Yes
26-24	PPPoE-Description	Client MAC address.	string: pppoe *client-mac-address*	No
26-25	Redirect-VRouter-Name	Client logical system:routing instance name indicating to which logical system:routing instance the request is redirected for user authentication.	string: *logical-system:routing-instance*	No
26-30	Tunnel-Nas-Port-Method	Method that determines whether the RADIUS server conveys to the LNS the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM. This information is conveyed only when the VSA value is 1. The VSA is formatted such that the first octet indicates the tunnel and the remaining three bytes are the attribute value.	4-octet integer: • 0 = none • 1 = Cisco CLID	Yes
26-31	Service-Bundle	SSC service bundle.	string *bundle-name*	No
26-33	Tunnel-Max-Sessions	Maximum number of sessions allowed in a tunnel.	integer: 4-octet	No
26-34	Framed-IP-Route-Tag	Route tag to apply to returned framed-ip-address.	integer: 4-octet	No
26-42	Input-Gigapackets	Number of times the input-packets attribute rolls over its 4-octet field.	integer	No
26-43	Output-Gigapackets	Number of times the output-packets attribute rolls over its 4-octet field.	integer	No
26-47	Ipv6-Primary-DNS	Client primary IPv6 DNS address negotiated by DHCP.	hexadecimal string: *ipv6-primary-dns-address*	No
26-48	Ipv6-Secondary-DNS	Client secondary IPv6 DNS address negotiated by DHCP.	hexadecimal string: *ipv6-secondary-dns-address*	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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 disconnected, and L2TP layer of the LNS initiates the termination. The PPP Disconnect Cause Code (L2TP AVP 46) is included in VSA 26-51 in the Accounting-Stop message that the router sends to the RADIUS server.	hexadecimal string: *disconnect-cause*	No
26-55	DHCP-Options	Client DHCP options.	string: *dhcp-options*	No
26-56	DHCP-MAC-Address	Client MAC address.	string: *mac-address*	No
26-57	DHCP-GI-Address	DHCP relay agent IP address.	integer: 4-octet	No
26-58	LI-Action	Traffic mirroring action. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs. If the CoA action is to stop mirroring (VSA 26-58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.	salt-encrypted integer 0=stop mirroring 1=start mirroring 2=no action	Yes
26-59	Med-Dev-Handle	Identifier that associates mirrored traffic to a specific subscriber. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.	salt-encrypted string	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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 which mirrored traffic is forwarded. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.	salt-encrypted IP address	No
26-61	Med-Port-Number	UDP port in the content destination device to which mirrored traffic is forwarded. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.	salt-encrypted integer	No
26-63	Interface-Desc	Text string that identifies the subscriber's access interface.	string: *interface-description*	No
26-64	Tunnel-Group	Name of the tunnel group (profile) assigned to a domain map.	string: *tunnel-group-name*	No
26-65	Activate-Service	Service to activate for the subscriber. Tagged VSA, which supports 8 tags (1-8).	string: *service-name*	Yes
26-66	Deactivate-Service	Service to deactivate for the subscriber.	string: *service-name*	Yes
26-67	Service-Volume	Amount of traffic, in MB, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8).	• range = 0 through 16777215 MB • 0 = no limit	Yes
26-68	Service-Timeout	Number of seconds that the service can be active; service is deactivated when the timeout expires. Tagged VSA, which supports 8 tags (1-8).	• range = 0 through 16777215 seconds • 0 = no timeout	Yes
26-69	Service-Statistics	Whether statistics for the service is enabled or disabled. Tagged VSA, which supports 8 tags (1-8).	• 0 = disable • 1 = enable time statistics • 2 = enable time and volume statistics	Yes
26-71	IGMP-Access-Name	Access list to use for the group (G) filter.	string: 32-octet	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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) filter.	string: 32-octet	Yes
26-74	MLD-Access-Name	Access list to use for the group (G) filter.	string: 32-octet	Yes
26-75	MLD-Access-Src-Name	Access list to use for the source-group (S,G) filter.	string: 32-octet	Yes
26-77	MLD-Version	MLD protocol version.	integer: 1-octet • 1=MLD version 1 • 2=MLD version 2	Yes
26-78	IGMP-Version	IGMP protocol version.	integer: 1-octet • 1=IGMP version 1 • 2=IGMP version 2 • 3=IGMP version 3	Yes
26-83	Service-Session	Name of the service.	string: *service-name*	No
26-84	Mobile-IP-Algorithm	Authentication algorithm used for Mobile IP registration.	integer: 4-octet	No
26-85	Mobile-IP-SPI	Security parameter index number for Mobile IP registration.	integer: 4-octet	No
26-86	Mobile-IP-Key	Security association MD5 key for Mobile IP registration.	string: key	No
26-87	Mobile-IP-Replay	Replay timestamp for Mobile IP registration.	integer: 4-octet	No
26-89	Mobile-IP-Lifetime	Registration lifetime for Mobile IP registration.	integer: 4-octet	No
26-91	Tunnel-Switch-Profile	Tunnel switch profile that determines whether a subscriber session is switched to a second session to a remote LNS. Takes precedence over tunnel switch profiles applied in any other manner.	string: *profile-name*	No
26-92	L2C-Up-Stream-Data	Actual upstream rate access loop parameter (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration.	string: actual upstream rate access loop parameter (ASCII encoded)	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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 (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration.	string: actual downstream rate access loop parameter (ASCII encoded)	No
26-94	Tunnel-Tx-Speed-Method	Method that determines the source from which the transmit speed is derived. Overrides global configuration in the CLI.	integer: 4-octet • 0 = none • 1 = static Layer 2 • 2 = dynamic layer 2. This method is not supported; the static Layer 2 method is used instead. • 3 = CoS. This method is not supported; the actual method is used instead. • 4 = actual • 5 = ANCP • 6 = PPPoE IA tags	No
26-97	IGMP-Immediate-Leave	IGMP Immediate Leave.	integer: 4-octet • 0=disable • 1=enable	Yes
26-100	MLD-Immediate-Leave	MLD Immediate Leave.	integer: 4-octet • 0=disable • 1=enable	Yes
26-106	IPv6-Ingress-Policy-Name	Input policy name to apply to a user IPv6 interface.	string: *policy-name*	Yes
26-107	IPv6-Egress-Policy-Name	Output policy name to apply to a user IPv6 interface.	string: *policy-name*	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 4: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-108	CoS-Traffic-Control-Profile-Parameter-Type	CoS traffic-shaping parameter type and description: • **T01**: Scheduler-map name • **T02**: Shaping rate • **T03**: Guaranteed rate • **T04**: Delay-buffer rate • **T05**: Excess rate • **T06**Traffic-control profile • **T07**: Shaping mode • **T08**: Byte adjust • **T09**: Adjust minimum • **T10**: Excess-rate high • **T11**: Excess-rate low • **T12**: Shaping rate burst • **T13**: Guaranteed rate burst	Two parts, delimited by white space: • Parameter type • Parameter value Examples: • **T01 smap_basic** • **T02 50m** • **T03 1m** • **T04 2000** • **T05 200** • **T06 tcp-gold** • **T07 frame-mode** • **T08 50**	Yes
26-109	DHCP-Guided-Relay-Server	IP address of DHCP server that DHCP relay agent uses to forward the discover PDUs.	integer: 4-byte *ip-address*	No
26-110	Acc-Loop-Cir-Id	Identification of the subscriber node connection to the access node.	string: up to 63 ASCII characters	No
26-111	Acc-Aggr-Cir-Id-Bin	Unique identification of the DSL line.	integer: 8-octet	No
26-112	Acc-Aggr-Cir-Id-Asc	Identification of the uplink on the access node, as in the following examples: • Ethernet access aggregation—**ethernet *slot/port [:inner-vlan-id] [:outer-vlan-id]*** • ATM aggregation—**atm *slot/port:vpi.vci***	string: up to 63 ASCII characters	No
26-113	Act-Data-Rate-Up	Actual upstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	No
26-114	Act-Data-Rate-Dn	Actual downstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	No
26-115	Min-Data-Rate-Up	Minimum upstream data rate configured for the subscriber.	integer: 4-octet	No
26-116	Min-Data-Rate-Dn	Minimum downstream data rate configured for the subscriber.	integer: 4-octet	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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 subscriber can attain.	integer: 4-octet	No
26-118	Att-Data-Rate-Dn	Maximum downstream data rate that the subscriber can attain.	integer: 4-octet	No
26-119	Max-Data-Rate-Up	Maximum upstream data rate configured for the subscriber.	integer: 4-octet	No
26-120	Max-Data-Rate-Dn	Maximum downstream data rate configured for the subscriber.	integer: 4-octet	No
26-121	Min-LP-Data-Rate-Up	Minimum upstream data rate in low power state configured for the subscriber.	integer: 4-octet	No
26-122	Min-LP-Data-Rate-Dn	Minimum downstream data rate in low power state configured for the subscriber.	integer: 4-octet	No
26-123	Max-Interlv-Delay-Up	Maximum one-way upstream interleaving delay configured for the subscriber.	integer: 4-octet	No
26-124	Act-Interlv-Delay-Up	Subscriber's actual one-way upstream interleaving delay..	integer: 4-octet	No
26-125	Max-Interlv-Delay-Dn	Maximum one-way downstream interleaving delay configured for the subscriber.	integer: 4-octet	No
26-126	Act-Interlv-Delay-Dn	Subscriber's actual one-way downstream interleaving delay.	integer: 4-octet	No
26-127	DSL-Line-State	State of the DSL line.	integer: 4-octet • 1 = Show uptime • 2 = Idle • 3 = Silent	No
26-128	DSL-Type	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated.	integer: 4-octet	No
26-130	Qos-Set-Name	Interface set to apply to the dynamic profile.	string: *interface-set-name*	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: 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 for this service. Tagged VSA, which supports 8 tags (1-8).	• range = 600 through 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).	Yes
26-141	Downstream-Calculated-QoS-Rate	Calculated (adjusted) downstream QoS rate in Kbps as set by the ANCP configuration.	range = 1000 through 4,294,967,295	No
26-142	Upstream-Calculated-QoS-Rate	Calculated (adjusted) upstream QoS rate in Kbps as set by the ANCP configuration.	range = 1000 through 4,294,967,295	No
26-143	Max-Clients-Per-Interface	Maximum allowable client sessions per interface. For DHCP clients, this value is the maximum sessions per logical interface. For PPPoE clients, this value is the maximum sessions (PPPoE interfaces) per PPPoE underlying interface.	integer: 4-octet	No
26-146	CoS-Scheduler-Pmt-Type	CoS scheduler parameter type and description: • Null: CoS scheduler name • **T01**: CoS scheduler transmit rate • **T02**: CoS scheduler buffer size • **T03**: CoS scheduler priority • **T04**: CoS scheduler drop-profile low • **T05**: CoS scheduler drop-profile medium-low • **T06**: CoS scheduler drop-profile medium-high • **T07**: CoS scheduler drop-profile high • **T08**: CoS scheduler drop-profile any	Three parts, delimited by white space: • Scheduler name • Parameter type • Parameter value Examples: • **be_sched** • **be_sched T01 12m** • **be_sched T02 26**	Yes
26-151	IPv6-Acct-Input-Octets	IPv6 receive octets.	integer	No
26-152	IPv6-Acct-Output-Octets	IPv6 transmit octets.	integer	No

Copyright © 2017, Juniper Networks, Inc.

Table 4: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-153	IPv6-Acct-Input-Packets	IPv6 receive packets.	integer	No
26-154	IPv6-Acct-Output-Packets	IPv6 transmit packets.	integer	No
26-155	IPv6-Acct-Input-Gigawords	IPv6 receive gigawords.	integer	No
26-156	IPv6-Acct-Output-Gigawords	IPv6 transmit gigawords.	integer	No
26-158	PPPoE-Padn	Route add for PPPoE sessions	string	No
26-160	Vlan-Map-Id	Trunk VLAN tag corresponding to the core-facing trunk physical interface.	integer	No
26-161	IPv6-Delegated-Pool-Name	Address pool used to locally allocate a delegated prefix (IA_PD).	string	No
26-162	Tx-Connect-Speed	Indication of transmit speed of the user's connection.	string	No
26-163	Rx-Connect-Speed	Indication of receive speed of the user's connection.	string	No
26-173	Service-Activate-Type	Indication of service activation type. This is a tagged attribute.	integer: 4-octet • 1 = dynamic-profile for residential services • 2 = op-script for business services	No
26-174	Client-Profile-Name	Enables RADIUS to override an assigned client dynamic profile with the included profile.	string	No
26-177	Cos-Shaping-Rate	Effective downstream shaping rate for subscriber.	string	No
26-179	Service-Volume-Gigawords	Amount of traffic, in 4GB units, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8).	• range = 0 through 16777215 4GB units • 0 = no limit	Yes
26-180	Update-Service	New values of service and time quotas for existing service. Tagged VSA, which supports 8 tags (1-8).	string: *service-name*	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 4: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-181	DHCPv6-Guided-Relay-Server	IPv6 addresses of DHCPv6 servers to which DHCPv6 relay agent forwards the Solicit and subsequent PDUs. Use multiple instances of the VSA to specify a list of servers.	hexadecimal string: *ipv6-address*	No
26-182	Acc-Loop-Remote-Id	Reports the ANCP Access-Loop-Remote-ID attribute.	string	No
26-183	Acc-Loop-Encap	Reports the ANCP Access-Loop-Encapsulation attribute.	hexadecimal string	No
26-184	Inner-Vlan-Map-Id	Inner VLAN tag allocated from the ranges provisioned on the core-facing physical interface, used to swap (replace) the autosensed VLAN tag on the access interface.	integer	No
26-185	Core-Facing-Interface	Name of the core-facing physical interface that forwards the Layer 2 wholesale session's downstream and upstream traffic relative to the network service provider (NSP) router.	string	No
26-187	DHCP-First-Relay-IPv4-Address	IPv4 address of the first relay link of a client/server binding.	integer: 4-byte *ip-address*	No
26-188	DHCP-First-Relay-IPv6-Address	IPv6 address of the first relay link of a client/server binding.	hexadecimal string: *ipv6-address*	No
26-191	Input-Interface-Filter	Name of an input filter to be attached to a **family any** interface.	string	No
26-192	Output-Interface-Filter	Name of an output filter to be attached to a **family any** interface.	string	No
26-194	Bulk-CoA-Transaction-Id	A common identifier or tag to associate the series of related CoA Requests as a transaction. This attribute is untagged and value 0 is reserved.	integer: 4-octet	Yes
26-195	Bulk-CoA-Identifier	A unique identifier for each CoA Request message that is part of the same transaction as specified by the Bulk-CoA-Transaction-Id VSA. This attribute is untagged and the value 0 is reserved.	integer: 4-octet	Yes

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 33
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39

AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 5 on page 33 shows the RADIUS attributes and Juniper Networks VSAs support in AAA access messages. A checkmark in a column indicates that the message type supports that attribute.

Table 5: 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	–	✓	✓	✓	–	–
22	Framed-Route	–	✓	–	–	–	–
25	Class	–	✓	–	–	–	–
26-1	Virtual-Router	✓	✓	–	–	–	–
26-4	Primary-DNS	–	✓	–	–	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 5: 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
26-5	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	✓	–	–	–	–	–
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	✓	–	–	–	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 5: 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	–	✓	–	–	–	–
26-94	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	✓	–	–	–	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 5: 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	✓	–	–	–	–	–
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	✓	–	–	–	–	–
26-143	Max-Clients-Per-Interface	–	✓	–	–	–	–
26-146	Cos-Scheduler-Pmt-Type	–	✓	–	–	✓	–
26-158	PPPoE-Padn	–	✓	–	–	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 5: 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
26-160	Vlan-Map-ID	–	✓	–	–	–	–
26-161	IPv6-Delegated-Pool-Name	–	✓	–	–	–	–
26-162	Tx-Connect-Speed	✓	–	–	–	–	–
26-163	Rx-Connect-Speed	✓	–	–	–	–	–
26-173	Service-Activate-Type	–	✓	–	–	✓	–
26-174	Client-Profile-Name	–	✓	–	–	–	–
26-179	Service-Volume-Gigawords	–	✓	–	–	✓	–
26-180	Update-Service	–	–	–	–	✓	–
26-181	DHCPv6-Guided-Relay-Server	–	✓	–	–	–	–
26-182	Acc-Loop-Remote-Id	✓	–	–	–	–	–
26-183	Acc-Loop-Encap	✓	–	–	–	–	–
26-184	Inner-Vlan-Map-Id	–	✓	–	–	–	–
26-185	Core-Facing-Interface	–	–	–	–	–	–
26-187	DHCP-First-Relay-IPv4-Address	✓	–	–	–	–	–
26-188	DHCP-First-Relay-IPv6-Address	✓	–	–	–	–	–
26-191	Input-Interface-Filter	✓	–	–	–	–	–
26-192	Output-Interface-Filter	✓	–	–	–	–	–
27	Session-Timeout	–	✓	–	✓	–	–
31	Calling-Station-ID	✓	–	–	–	✓	–
32	NAS-Identifier	✓	–	–	–	–	–
44	Acct-Session-ID	✓	–	–	–	✓	✓
61	NAS-Port-Type	✓	–	–	–	–	–
64	Tunnel-Type	✓	✓	–	–	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 5: 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
65	Tunnel-Medium-Type	✓	✓	–	–	–	–
66	Tunnel-Client-Endpoint	✓	✓	–	–	–	–
67	Tunnel-Server-Endpoint	✓	✓	–	–	–	–
68	Acct-Tunnel-Connection-ID	✓	✓	–	–	–	–
68	Acct-Tunnel-Connection	–	✓	–	–	–	–
69	Tunnel-Password	–	✓	–	–	–	–
82	Tunnel-Assignment-Id	✓	✓	–	–	–	–
83	Tunnel-Preference	–	✓	–	–	–	–
85	Acct-Interim-Interval	–	✓	–	–	–	–
87	NAS-Port-Id	✓	–	–	–	✓	–
88	Framed-Pool	–	✓	–	–	–	–
90	Tunnel-Client-Auth-Id	✓	✓	–	–	–	–
91	Tunnel-Server-Auth-Id	✓	✓	–	–	–	–
96	Framed-Interface-ID	–	✓	–	–	–	–
97	Framed-IPv6-Prefix	–	✓	–	–	–	–
99	Framed-IPv6-Route	–	✓	–	–	–	–
100	Framed-IPv6-Pool	–	✓	–	–	–	–
101	Error-Cause	–	–	–	–	✓	✓
123	Delegated-IPv6-Prefix	–	✓	–	–	–	–
242	Ascend-Data-Filter	–	✓	–	–	✓	–

Related Documentation

- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 13

Copyright © 2017, Juniper Networks, Inc.

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 14
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21

AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 6 on page 39 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 6: 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	✓	✓	✓	–	–
26-42	Input-Gigapackets	–	✓	✓	–	–
26-43	Output-Gigapackets	–	✓	✓	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 6: 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
26-47	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	–	✓	✓	–	–
26-110	Acc-Loop-Cir-Id	✓	✓	✓	–	–
26-111	Acc-Aggr-Cir-Id-Bin	✓	✓	✓	–	–
26-112	Acc-Aggr-Cir-Id-Asc	✓	✓	✓	–	–
26-113	Act-Data-Rate-Up	✓	✓	✓	–	–
26-114	Act-Data-Rate-Dn	✓	✓	✓	–	–
26-115	Min-Data-Rate-Up	✓	✓	✓	–	–
26-116	Min-Data-Rate-Dn	✓	✓	✓	–	–
26-117	Att-Data-Rate-Up	✓	✓	✓	–	–
26-118	Att-Data-Rate-Dn	✓	✓	✓	–	–
26-119	Max-Data-Rate-Up	✓	✓	✓	–	–
26-120	Max-Data-Rate-Dn	✓	✓	✓	–	–
26-121	Min-LP-Data-Rate-Up	✓	✓	✓	–	–
26-122	Min-LP-Data-Rate-Dn	✓	✓	✓	–	–
26-123	Max-Interlv-Delay-Up	✓	✓	✓	–	–
26-124	Act-Interlv-Delay-Up	✓	✓	✓	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 6: 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
26-125	Max-Interlv-Delay-Dn	✓	✓	✓	–	–
26-126	Act-Interlv-Delay-Dn	✓	✓	✓	–	–
26-127	DSL-Line-State	✓	✓	✓	–	–
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	✓	✓	✓	–	–
26-177	Cos-Shaping-Rate	✓	✓	✓	–	–
26-182	Acc-Loop-Remote-Id	✓	✓	–	–	–
26-183	Acc-Loop-Encap	✓	✓	–	–	–
26-184	Inner-Vlan-Map-Id	✓	✓	–	–	–
26-185	Core-Facing-Interface	✓	✓	–	–	–
26-187	DHCP-First-Relay-IPv4-Address	✓	✓	✓	–	–
26-188	DHCP-First-Relay-IPv6-Address	✓	✓	✓	–	–
26-191	Input-Interface-Filter	✓	✓	✓	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 6: 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
26-192	Output-Interface-Filter	✓	✓	✓	–	–
31	Calling-Station-ID	✓	✓	✓	–	–
32	NAS-Identifier	✓	✓	✓	–	–
40	Acct-Status-Type	✓	✓	✓	✓	✓
41	Acct-Delay-Time	✓	✓	✓	✓	✓
42	Acct-Input-Octets	–	✓	✓	–	–
43	Acct-Output-Octets	–	✓	✓	–	–
44	Acct-Session-ID	✓	✓	✓	✓	✓
45	Acct-Authentic	✓	✓	✓	✓	✓
46	Acct-Session-Time	–	✓	✓	–	–
47	Acct-Input-Packets	–	✓	✓	–	–
48	Acct-Output-Packets	–	✓	✓	–	–
49	Acct-Terminate-Cause	–	✓	✓	–	–
52	Acct-Input-Gigawords	–	✓	✓	–	–
53	Acct-Output-Gigawords	–	✓	✓	–	–
55	Event-Timestamp	✓	✓	✓	✓	✓
61	NAS-Port-Type	✓	✓	✓	–	–
64	Tunnel-Type	✓	✓	✓	–	–
65	Tunnel-Medium-Type	✓	✓	✓	–	–
66	Tunnel-Client-Endpoint	✓	✓	✓	–	–
67	Tunnel-Server-Endpoint	✓	✓	✓	–	–
68	Acct-Tunnel-Connection	✓	✓	✓	–	–
82	Tunnel-Assignment-Id	✓	✓	✓	–	–

Copyright © 2017, Juniper Networks, Inc.

Table 6: 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
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	✓	✓	✓	–	–

Related Documentation

- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 33
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 13
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 14
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21

Configuring How RADIUS Attributes Are Used for Subscriber Access

You can specify the attributes RADIUS ignores in RADIUS Access-Accept messages, and the attributes RADIUS excludes from specified message types.

To configure the attributes RADIUS ignores or excludes:

1. Specify that you want to configure RADIUS.

   ```
   [edit access profile isp-bos-metro-fiber-basic]
   user@host# edit radius
   ```

2. Specify that you want to configure how RADIUS attributes are ignored or excluded.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius]
   user@host# edit attributes
   ```

3. Specify the attributes you want RADIUS to ignore when the attributes are in Access-Accept messages. See Table 7 on page 44 for the attributes you can configure.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius attributes]
   user@host# set ignore input-filter output-filter
   ```

Copyright © 2017, Juniper Networks, Inc.

4. Configure RADIUS to exclude the specified attribute from the specified RADIUS message type. See Table 8 on page 44 for the attributes and message type combinations you can configure.

   [edit access profile isp-bos-metro-fiber-basic radius attributes]
   user@host# set exclude input-filter output-filter

You use the **ignore** statement to configure the router or switch to ignore a particular attribute in RADIUS Access-Accept messages. By default, the router or switch processes the attributes received from the external AAA server. Table 7 on page 44 lists the attributes supported in the **ignore** statement.

Table 7: Attributes That Can Be Ignored in RADIUS Access-Accept Messages

CLI Entry	Attribute Name	Attribute Number
dynamic-iflset-name	Interface-Set-Name	Juniper Networks VSA 26-130
framed-ip-netmask	Framed-Ip-Netmask	RADIUS attribute 9
input-filter	Ingress-Policy-Name	Juniper Networks VSA 26–10
logical-system:routing-instance	Virtual-Router	Juniper Networks VSA 26–1
output-filter	Egress-Policy-Name	Juniper Networks VSA 26–11

You use the **exclude** statement to configure the router or switch to exclude the specified attributes from the specified type of RADIUS message. Not all attributes appear in all types of RADIUS messages—the CLI indicates the RADIUS message type. By default, the router or switch includes the specified attributes in RADIUS Access-Request, Accounting-On, Accounting-Off, Accounting-Start, Accounting-Stop, and Interim-Accounting messages. Table 8 on page 44 lists the attributes and message types supported in the **exclude** statement.

Table 8: Attributes That Can Be Excluded from RADIUS Messages

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
accounting-authentic	Acct-Authentic	RADIUS attribute 45	Accounting-On Accounting-Off
accounting-delay-time	Acct-Delay-Time	RADIUS attribute 41	Accounting-On Accounting-Off

Copyright © 2017, Juniper Networks, Inc.

Table 8: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
accounting-session-id	Acct-Session-Id	RADIUS attribute 44	Access-Request Accounting-On Accounting-Off Accounting-Stop
accounting-terminate-cause	Acct-Terminate-Cause	RADIUS attribute 49	Accounting-Off
called-station-id	Called-Station-Id	RADIUS attribute 30	Access-Request Accounting-Start Accounting-Stop
calling-station-id	Calling-Station-Id	RADIUS attribute 31	Access-Request Accounting-Start Accounting-Stop
class	Class	RADIUS attribute 25	Accounting-Start Accounting-Stop
cos-shaping-rate	Cos-Shaping-Rate	Juniper Networks VSA 26-177	Accounting-Start Accounting-Stop
delegated-ipv6-prefix	Delegated-IPv6-Prefix	RADIUS attribute 123	Accounting-Start Accounting-Stop
dhcp-gi-address	DHCP-GI-Address	Juniper Networks VSA 26–57	Access-Request Accounting-Start Accounting-Stop
dhcp-mac-address	DHCP-MAC-Address	Juniper Networks VSA 26–56	Access-Request Accounting-Start Accounting-Stop
dhcp-options	DHCP-Options	Juniper Networks VSA 26–55	Access-Request Accounting-Start Accounting-Stop

Copyright © 2017, Juniper Networks, Inc.

Table 8: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
downstream-calculated-qos-rate	Downstream-Calculated-QoS-Rate	Juniper Networks VSA 26-141	Access-Request Accounting-Start Accounting-Stop Interim-Accounting
dsl-forum-attributes	Not applicable	Excludes the DSL Forum VSA (IANA vendor ID 3561)	Access-Request Accounting-Start Accounting-Stop Interim-Accounting
dynamic-iflset-name	Qos-Set-Name	Juniper Networks VSA 26–130	Accounting-Start Accounting-Stop
event-timestamp	Event-Timestamp	RADIUS attribute 55	Accounting-On Accounting-Off Accounting-Start Accounting-Stop
filter-id	Filter-Id	RADIUS attribute 11	Accounting-Start Accounting-Stop
first-relay-ipv4-address	DHCP-First-Relay-IPv4-Address	Juniper Networks VSA 26–187	Accounting-Start Accounting-Stop Interim-Accounting
first-relay-ipv6-address	DHCP-First-Relay-IPv6-Address	Juniper Networks VSA 26–188	Accounting-Start Accounting-Stop Interim-Accounting
framed-ip-address	Framed-IP-Address	RADIUS attribute 8	Accounting-Start Accounting-Stop
framed-ip-netmask	Framed-IP-Netmask	RADIUS attribute 9	Accounting-Start Accounting-Stop

Copyright © 2017, Juniper Networks, Inc.

Table 8: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
framed-ip-route	Framed-Route	RADIUS attribute 22	Accounting-Start Accounting-Stop
framed-ipv6-pool	Framed-IPv6-Pool	RADIUS attribute 100	Accounting-Start Accounting-Stop
framed-ipv6-prefix	Framed-IPv6-Prefix	RADIUS attribute 97	Accounting-Start Accounting-Stop
framed-ipv6-route	Framed-IPv6-Route	RADIUS attribute 99	Accounting-Start Accounting-Stop
framed-pool	Framed-Pool	RADIUS attribute 88	Accounting-Start Accounting-Stop
input-filter	Ingress-Policy-Name	Juniper Networks VSA 26–10	Accounting-Start Accounting-Stop
input-gigapackets	Acct-Input-Gigapackets	Juniper Networks VSA 26–42	Accounting-Stop
input-gigawords	Acct-Input-Gigawords	RADIUS attribute 52	Accounting-Stop
input-ipv6-gigawords	IPv6-Acct-Input-Gigawords	Juniper Networks VSA 26–155	Accounting-Stop
input-ipv6-octets	IPv6-Acct-Input-Octets	Juniper Networks VSA 26–151	Accounting-Stop
input-ipv6-packets	IPv6-Acct-Input-Packets	Juniper Networks VSA 26–153	Accounting-Stop
interface-description	Interface-Desc	Juniper Networks VSA 26–53	Access-Request Accounting-Start Accounting-Stop

Copyright © 2017, Juniper Networks, Inc.

Table 8: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
nas-identifier	NAS-Identifier	RADIUS attribute 32	Access-Request Accounting-On Accounting-Off Accounting-Start Accounting-Stop
nas-port	NAS-Port	RADIUS attribute 5	Access-Request Accounting-Start Accounting-Stop
nas-port-id	NAS-Port-Id	RADIUS attribute 87	Access-Request Accounting-Start Accounting-Stop
nas-port-type	NAS-Port-Type	RADIUS attribute 61	Access-Request Accounting-Start Accounting-Stop
output-filter	Egress-Policy-Name	Juniper Networks VSA 26–11	Accounting-Start Accounting-Stop
pppoe-description	PPPoE Description	Juniper Networks VSA 26-24	Access-Request Accounting-Start Accounting-Stop Interim-Accounting
output-gigapackets	Acct-Output-Gigapackets	Juniper Networks VSA 26–43	Accounting-Stop
output-gigawords	Acct-Output-Gigawords	RADIUS attribute 53	Accounting-Stop
output-ipv6-gigawords	IPv6-Acct-Output-Gigawords	Juniper Networks VSA 26–156	Accounting-Stop
output-ipv6-octets	IPv6-Acct-Output-Octets	Juniper Networks VSA 26–152	Accounting-Stop

Copyright © 2017, Juniper Networks, Inc.

Table 8: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
output-ipv6-packets	IPv6-Acct-Output-Packets	Juniper Networks VSA 26–154	Accounting-Stop
upstream-calculated-qos-rate	Upstream-Calculated-QoS-Rate	Juniper Networks VSA 26-142	Access-Request Accounting-Start Accounting-Stop Interim-Accounting
virtual-router	Virtual-Router	Juniper Networks VSA 26-1	Access-Request Accounting-Start Accounting-Stop

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 93
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs

Table 9 on page 50 lists the RADIUS attributes and Juniper Networks VSAs and their corresponding Junos OS predefined variables that are used in dynamic profiles. When the router instantiates a dynamic profile following subscriber access, the Junos OS uses the predefined variable to specify the RADIUS attribute or VSA for the information obtained from the RADIUS server.

Some predefined variables support the configuration of default values. The configured default value is used in the event that RADIUS fails to return a value for the variable. You configure default values with the **predefined-variable-defaults *predefined-variable default-value*** statement at the **[edit dynamic-profiles]** hierarchy level. When you specify the ***predefined-variable***, you use the name of the Junos OS predefined variable, but you omit the leading **$junos-** prefix.

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
RADIUS Attribute			
Framed-IP-Address (8)	$junos-framed-route-ip-address	Address for the client	No
Filter-ID (11)	$junos-input-filter NOTE: Variable is also used for VSA 26–10.	Input filter to apply to client IPv4 interface	Yes
Framed-Route (22)	$junos-framed-route-ip-address-prefix	(Subattribute 1): Route prefix for access route	No
	$junos-framed-route-nexthop	(Subattribute 2): Next hop address for access route	No
	$junos-framed-route-cost	(Subattribute 3): Metric for access route	No
	$junos-framed-route-distance	(Subattribute 5): Preference for access route	No
	$junos-framed-route-tag	(Subattribute 6): Tag for access route	No
Framed-IPv6-Prefix (97)	$junos-ipv6-ndra-prefix	Prefix value in IPv6 Neighbor Discovery route advertisements	No
Framed-IPv6-Route (99)	$junos-framed-route-ipv6-address-prefix	(Subattribute 1): Framed IPv6 route prefix configured for the client	No
	$junos-framed-route-ipv6-cost	(Subattribute 3): Metric for access route	No
	$junos-framed-route-ipv6-distance	(Subattribute 5): Preference for access route	No
	$junos-framed-route-ipv6-nexthop	(Subattribute 2): IPv6 routing information configured for the client	No
	$junos-framed-route-ipv6-tag	(Subattribute 6): Tag for access route	No
Juniper Networks VSA			

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
Virtual-Router (26–1)	$junos-routing-instance	Routing instance to which subscriber is assigned	Yes
Ingress-Policy-Name (26–10)	$junos-input-filter NOTE: Variable is also used for RADIUS attribute 11.	Input filter to apply to client IPv4 interface	Yes
Egress-Policy-Name (26–11)	$junos-output-filter	Output filter to apply to client IPv4 interface	Yes
IGMP-Enable (26–23)	$junos-igmp-enable	Enable or disable IGMP on client interface	Yes
IGMP-Access-Name (26–71)	$junos-igmp-access-group-name	Access list to use for the group (G) filter	Yes
IGMP-Access-Src-Name (26–72)	$junos-igmp-access-source-group-name	Access List to use for the source group (S,G) filter	Yes
MLD-Access-Name (26–74)	$junos-mld-access-group-name	Access list to use for the group (G) filter	Yes
MLD-Access-Src-Name (26–75)	$junos-mld-access-source-group-name	Access List to use for the source group (S,G) filter	Yes
MLD-Version (26–77)	$junos-mld-version	MLD protocol version	Yes
IGMP-Version (26–78)	$junos-igmp-version	IGMP protocol version	Yes
IGMP-Immediate-Leave (26–97)	$junos-igmp-immediate-leave	IGMP immediate leave	Yes
MLD-Immediate-Leave (26–100)	$junos-mld-immediate-leave	MLD immediate leave	Yes
IPv6-Ingress-Policy-Name (26–106)	$junos-input-ipv6-filter	Input filter to apply to client IPv6 interface	Yes
IPv6-Egress-Policy-Name (26–107)	$junos-output-ipv6-filter	Output filter to apply to client IPv6 interface	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
CoS-Traffic-Control-Profile-Parameter-Type (26–108)	$junos-cos-scheduler-map	(T01: Scheduler-map name) Name of scheduler map configured in traffic-control profile	Yes
	$junos-cos-shaping-rate	(T02: Shaping rate) Shaping rate configured in traffic-control profile	Yes
	$junos-cos-guaranteed-rate	(T03: Guaranteed rate) Guaranteed rate configured in traffic-control profile	Yes
	$junos-cos-delay-buffer-rate	(T04: Delay-buffer rate) Delay-buffer rate configured in traffic-control profile	Yes
	$junos-cos-excess-rate	(T05; Excess rate) Excess rate configured in traffic-control profile	Yes
	$junos-cos-traffic-control-profile	(T06: Traffic-control profile) Name of the traffic-control profile configured in a dynamic profile	Yes
	$junos-cos-shaping-mode	(T07; Shaping mode) CoS shaping mode configured in a dynamic profile	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
	$junos-cos-byte-adjust	(T08; Byte adjust) Byte adjustments configured for the shaping mode in a dynamic profile	Yes
	$junos-cos-adjust-minimum	(T09; Adjust minimum) Minimum adjusted value allowed for the shaping rate in a dynamic profile	Yes
	$junos-cos-excess-rate-high	(T10; Excess rate high) Excess rate configured for high-priority traffic in a dynamic profile	Yes
	$junos-cos-excess-rate-low	(T11; Excess rate low) Excess rate configured for low-priority traffic in a dynamic profile	Yes
	$junos-cos-shaping-rate-burst	(T12; Shaping rate burst) Burst size configured for the shaping rate in a dynamic profile	Yes
	$junos-cos-guaranteed-rate-burst	(T13; Guaranteed rate burst) Burst size configured for the guaranteed rate in a dynamic profile	Yes
Qos-Set-Name (26–130)	$junos-interface-set-name	Name of an interface set configured in a dynamic profile	Yes
CoS-Scheduler-Pmt-Type (26–146)	$junos-cos-scheduler	(Null: Scheduler name) Name of scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-tx	(T01: CoS scheduler transmit rate) Transmit rate for scheduler configured in a dynamic profile	Yes Available for multiple parameters: • Percent • Rate

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
	$junos-cos-scheduler-bs	(T02: CoS scheduler buffer size) Buffer size for scheduler configured in a dynamic profile	Yes Available for multiple parameters: • Percent • Temporal
	$junos-cos-scheduler-pri	(T03: CoS scheduler priority) Packet-scheduling priority for scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-dropfile-low	(T04: CoS scheduler drop-profile low) Name of drop profile for RED loss-priority level **low** for scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-dropfile-medium-low	(T05: CoS scheduler drop-profile medium-low) Name of drop profile for RED loss-priority level **medium-low** for scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-dropfile-medium-high	(T06: CoS scheduler drop-profile medium-high) Name of drop profile for RED loss-priority level **medium-high** for scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-dropfile-high	(T07: CoS scheduler drop-profile high) Name of drop profile for RED loss-priority level **high** for scheduler configured in a dynamic profile	Yes

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
	$junos-cos-scheduler-dropfile-any	(T08: CoS scheduler drop-profile any) Name of drop profile for RED loss-priority level **any** for scheduler configured in a dynamic profile	Yes
	$junos-cos-scheduler-excess-rate	(T09: CoS scheduler excess rate) Excess rate configured for a scheduler in a dynamic profile	Yes Available for multiple parameters: • Percent • Proportion
	$junos-cos-scheduler-shaping-rate	(T10: CoS scheduler shaping rate) Shaping rate configured for a scheduler in a dynamic profile	Yes Available for multiple parameters: • Percent • Rate
	$junos-cos-scheduler-excess-priority	(T11: CoS scheduler excess priority) Excess priority configured for a scheduler in a dynamic profile	Yes
Input-Interface-Filter (26-191)	$junos-input-interface-filter	Name of an input filter to be attached to a **family any** interface.	No
Output-Interface-Filter (26-192)	$junos-output-interface-filter	Name of an output filter to be attached to a **family any** interface.	No
IPv4 input service set (26-196)	$junos-input-service-set	Name of an IPv4 input service set to be attached.	No
IPv4 output service set (26-197)	$junos-output-service-set	Name of an IPv4 output service set to be attached.	No
IPv4 input service filter (26-198)	$junos-input-service-filter	Name of an IPv4 input service filter to be attached.	No

Copyright © 2017, Juniper Networks, Inc.

Table 9: RADIUS Attributes and Corresponding Junos OS Predefined Variables *(continued)*

RADIUS Attribute or VSA	Junos OS Predefined Variable	Description	Default Value Support for Junos OS Predefined Variable
IPv4 output service filter (26-199)	$junos-output-service-filter	Name of an IPv4 output service filter to be attached.	No
IPv6 input service set (26-200)	$junos-input-ipv6-service-set	Name of an IPv6 input service set to be attached.	No
IPv6 output service set (26-201)	$junos-output-ipv6-service-set	Name of an IPv6 output service set to be attached.	No
IPv6 input service filter (26-202)	$junos-input-ipv6-service-filter	Name of an IPv6 input service filter to be attached.	No
IPv6 output service filter (26-203)	$junos-output-ipv6-service-filter	Name of an IPv6 output service filter to be attached.	No
PCEF profile (26-204)	$junos-pcef-profile	Name of a PCEF profile to be attached.	No
PCEF rule (26-205)	$junos-pcef-rule	Name of a PCC rule to activate.	No

Related Documentation

- *Dynamic Variables Overview*
- *Configuring Predefined Dynamic Variables in Dynamic Profiles*
- *Junos OS Predefined Variables*

DSL Forum Vendor-Specific Attributes

Digital Subscriber Line (DSL) attributes are RADIUS vendor-specific attributes (VSAs) that are defined by the DSL Forum. The attributes transport DSL information that is not supported by standard RADIUS attributes and which convey information about the associated DSL subscriber and data rate. The attributes are defined in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*.

NOTE: **Junos OS uses the vendor ID 3561, which is assigned by the Internet Assigned Numbers Authority (IANA), for the DSL Forum VSAs.**

Subscriber management does not process DSL values—the router simply passes the values received from the subscriber to the RADIUS server, without performing any parsing

Copyright © 2017, Juniper Networks, Inc.

or manipulation. However, you can manage the content of DSL VSA values either by using the client configuration to restrict the DSL VSAs that the client sends, or by configuring the RADIUS server to ignore unwanted DSL VSAs.

Table 10 on page 57 describes the DSL Forum VSAs.

Table 10: DSL Forum VSAs

Attribute Number	Attribute Name	Description	Value
[26-1]	Agent-Circuit-Id	Identifier for the subscriber agent circuit ID (ACI) that corresponds to the DSLAM 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.	string
[26-2]	Agent-Remote-Id	Unique identifier for the subscriber associated with the DSLAM 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.	string
[26-129]	Actual-Data-Rate-Upstream	Actual upstream data rate of the subscriber's synchronized DSL link	integer: 4-octet
[26-130]	Actual-Data-Rate-Downstream	Actual downstream data rate of the subscriber's synchronized DSL link	integer: 4-octet
[26-131]	Minimum-Data-Rate-Upstream	Minimum upstream data rate configured for the subscriber	integer: 4-octet
[26-132]	Minimum-Data-Rate-Downstream	Minimum downstream data rate configured for the subscriber	integer: 4-octet
[26-133]	Attainable-Data-Rate-Upstream	Upstream data rate that the subscriber can attain	integer: 4-octet
[26-134]	Attainable-Data-Rate-Downstream	Downstream data rate that the subscriber can attain	integer: 4-octet
[26-135]	Maximum-Data-Rate-Upstream	Maximum upstream data rate configured for the subscriber	integer: 4-octet

Copyright © 2017, Juniper Networks, Inc.

Table 10: DSL Forum VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value
[26-136]	Maximum-Data-Rate-Downstream	Maximum downstream data rate configured for the subscriber	integer: 4-octet
[26-137]	Minimum-Data-Rate-Upstream-Low-Power	Minimum upstream data rate in low power state configured for the subscriber	integer: 4-octet
[26-138]	Minimum-Data-Rate-Downstream-Low-Power	Minimum downstream data rate in low power state configured for the subscriber	integer: 4-octet
[26-139]	Maximum-Interleaving-Delay-Upstream	Maximum one-way upstream interleaving delay configured for the subscriber	integer: 4-octet
[26-140]	Actual-Interleaving-Delay-Upstream	Subscriber's actual one-way upstream interleaving delay	integer: 4-octet
[26-141]	Maximum-Interleaving-Delay-Downstream	Maximum one-way downstream interleaving delay configured for the subscriber	integer: 4-octet
[26-142]	Actual-Interleaving-Delay-Downstream	Subscriber's actual one-way downstream interleaving delay	integer: 4-octet
[26-144]	Access-Loop-Encapsulation	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated	string: 3-byte
[26-254]	IWF-Session	Indication that the interworking function (IWF) has been performed for the subscriber's session	No data field required

Related Documentation

- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS on page 58

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

Table 11 on page 58 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. The DSL Forum vendor ID is 3561 (hexadecimal DE9), which is assigned by the IANA.

Table 11: DSL Forum VSAs—Supported RADIUS Messages

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
[26-1]	Agent-Circuit-Id	✓	✓	✓	✓	✓

Copyright © 2017, Juniper Networks, Inc.

Table 11: DSL Forum VSAs—Supported RADIUS Messages *(continued)*

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
[26-2]	Agent-Remote-Id	✓	✓	✓	✓	✓
[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	✓	✓	✓	✓	–
[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	✓	✓	✓	✓	–
[26-254]	IWF-Session	✓	✓	✓	✓	–

Related Documentation

- DSL Forum Vendor-Specific Attributes on page 56

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.

Copyright © 2017, Juniper Networks, Inc.

The two VSAs are shown in the following table, and are described in RFC 2548 (*Microsoft Vendor-specific RADIUS Attributes*)

Table 12: 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*

Release History Table

Release	Description
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

- DNS Address Assignment Precedence on page 256

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 3

Configuring RADIUS NAS-Port Attributes and Options

- Manual Configuration of the NAS-Port-ID RADIUS Attribute on page 61
- Configuring a NAS-Port-ID with Additional Options on page 63
- Configuring the Order in Which Optional Values Appear in the NAS-Port-ID on page 64
- Configuring a Calling-Station-ID with Additional Attributes on page 66
- Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers on page 70
- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 71
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Manual Configuration of the NAS-Port-Type RADIUS Attribute on page 75
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85
- Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces on page 88

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

Copyright © 2017, Juniper Networks, Inc.

also specify an optional delimiter character, which separates the values in a NAS-Port-ID. The default delimiter character is the hash character (**#**).

A default NAS-Port-ID consists of the following **interface-description** string:

[physical-interface].<interface-type>-<slot>/<adapter>/<port><.subinterface>[:<svlan>-<vlan>]

For example: **ge-1/2/0.100:100**

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

Related Documentation

- Configuring a NAS-Port-ID with Additional Options on page 63
- RADIUS Server Options for Subscriber Access on page 4

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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 64.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring the Order in Which Optional Values Appear in the NAS-Port-ID on page 64

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring a NAS-Port-ID with Additional Options on page 63

Copyright © 2017, Juniper Networks, Inc.

Configuring a Calling-Station-ID with Additional Attributes

You can configure an alternative value for the Calling-Station-ID (RADIUS IETF attribute 31) in an access profile on the MX Series router.

By default, the Calling-Station-ID includes the **agent-circuit-id** string. Optionally, you can configure the Calling-Station-ID to include one or more of the following attributes, in any combination:

- 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 additional attributes 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]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# edit radius options
```

3. Specify the nondefault character to use as the delimiter between the different attribute 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.

Copyright © 2017, Juniper Networks, Inc.

[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 Attributes 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 optional attributes.

```
[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

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- RADIUS Server Options for Subscriber Access on page 4

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
     ```

Related Documentation

- unique-nas-port on page 1436

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

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 on page 72
- NAS-Port RADIUS Attribute on page 72
- NAS-Port Options Configuration and Subscriber Network Access Models on page 72
- NAS-Port Options Definition on page 72

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

- 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.

Related Documentation

- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74

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.

Related Documentation

- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 71
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74

Copyright © 2017, Juniper Networks, Inc.

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

Copyright © 2017, Juniper Networks, Inc.

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 77.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per VLAN" on page 79.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per Stacked VLAN" on page 80.

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 82.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per VLAN" on page 84.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN" on page 85.

Related Documentation

- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 71
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

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

Copyright © 2017, Juniper Networks, Inc.

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 13 on page 76 shows the supported port type values for RADIUS attribute 61 (NAS-Port-Type) that you can include in an access profile.

Table 13: 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)
idsl	14	ISDN DSL
isdn-sync	2	ISDN Synchronous
isdn-v110	4	ISDN Async V.110
isdn-v120	3	ISDN Async V.120

Copyright © 2017, Juniper Networks, Inc.

Table 13: RADIUS NAS-Port-Type Values *(continued)*

Statement Option	NAS-Port-Type Value	Description
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
x75	9	X.75
xdsl	16	DSL of unknown type

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117

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]

Copyright © 2017, Juniper Networks, Inc.

```
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;
  }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79

Copyright © 2017, Juniper Networks, Inc.

- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

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***)

Copyright © 2017, Juniper Networks, Inc.

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;
  }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
    nas-port-type 4722;
    stacked-vlan-ranges {
      2-10,any;
    }
  }
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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;
  }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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;
  }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 85

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]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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.

Copyright © 2017, Juniper Networks, Inc.

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;
    }
  }
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 77
- Configuring the RADIUS NAS-Port-Type per VLAN on page 79
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 80
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 82
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 84

Copyright © 2017, Juniper Networks, Inc.

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]
```


Copyright © 2017, Juniper Networks, Inc.

```
nas-port-options boston-subscribers {
  nas-port-extended-format {
    slot-width 6;
    adapter-width 3;
    port-width 4;
    vpi-width 12;
    vci-width 24;
  }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- RADIUS Server Options for Subscriber Access on page 4
- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 4

Configuring RADIUS Authentication for Subscriber Access

- Retaining Authentication and Accounting Information During Session Startup on page 91
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Specifying the Authentication and Accounting Methods for Subscriber Access on page 92
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 93
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Retaining Authentication and Accounting Information During Session Startup

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 optionally use the **wait-for-acct-on-ack** statement to configure the **authd** process 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.

You configure this feature for an access profile for a logical system and routing instance context. All authentication requests fail until the router receives an Acct-On-Ack response from a RADIUS accounting server that is configured in the access profile. If multiple RADIUS accounting servers are configured for the access profile, **authd** waits until the first response is received.

You can also configure the **authd** process to send accounting messages when the RADIUS server status changes for an access profile. This configuration enables you to monitor

Copyright © 2017, Juniper Networks, Inc.

whether the access profile has an active RADIUS server. You use the **send-acct-status-on-config-change** statement to specify that **authd** send 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 deleted from the access profile.

Related Documentation

- Configuring Per-Subscriber Session Accounting on page 105

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 92.

2. Specify how accounting statistics are collected.

   See "Configuring Per-Subscriber Session Accounting" on page 105.

Related Documentation

- AAA Service Framework Overview on page 3
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

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 and, if it fails, local authentication (**password**) is done.

Copyright © 2017, Juniper Networks, Inc.

You can specify the following authentication methods:

NOTE: You must always specify the **radius** authentication method. Subscriber access management does not support the **password** keyword (the default), and authentication fails when no method is specified.

- **password**—Local authentication
- **radius**—RADIUS-based authentication

You can specify the following accounting methods:

- **radius**—RADIUS-based accounting

To configure the authentication and accounting methods for subscriber access management:

1. Specify the authentication methods and the order in which they are used. Only **radius** is supported.

   ```
   [edit access profile isp-bos-metro-fiber-basic]
   user@host# set authentication-order radius
   ```

2. Specify the accounting method.

   ```
   [edit access profile isp-bos-metro-fiber-basic]
   user@host# set accounting order radius
   ```

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring Per-Subscriber Session Accounting on page 105
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

Related Documentation

- Attaching Access Profiles on page 124
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Example: Configuring RADIUS-Based Subscriber Authentication and Accounting

This example shows a RADIUS-based authentication and accounting configuration.

```
[edit access]
radius-server {
    192.168.1.250 {
        port 1812;
        accounting-port 1813;
        accounting-retry 6;
        accounting-timeout 20;
        retry 3;
        secret $ABC123$ABC123;
        source-address 192.168.1.100;
        timeout 45;
    }
    192.168.1.251 {
        port 1812;
        accounting-port 1813;
        accounting-retry 6;
        accounting-timeout 20;
        retry 3;
        secret $ABC123;
        source-address 192.168.1.100;
```


Copyright © 2017, Juniper Networks, Inc.

```
      timeout 30;
    }
    2001:DB8:0f101::2{
      port 1812;
      accounting-port 1813;
      accounting-retry 6;
      accounting-timeout 20;
      retry 4;
      secret $ABC123$ABC123$ABC123-;
      source-address 2001:DB8:0f101::1;
      timeout 20;
    }
  }
  profile isp-bos-metro-fiber-basic {
    authentication-order radius;
    accounting {
      order radius;
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      immediate-update;
      statistics time;
      update-interval 12;
      wait-for-acct-on-ack;
      send-acct-status-on-config-change;
    }
    radius {
      authentication-server 192.168.1.251 192.168.1.252;
      accounting-server 192.168.1.250 192.168.1.251;
      options {
        accounting-session-id-format decimal;
        client-accounting-algorithm round-robin;
        client-authentication-algorithm round-robin;
        nas-identifier 56;
        nas-port-id-delimiter %;
        nas-port-id-format {
          nas-identifier;
          interface-description;
        }
        nas-port-type {
          ethernet {
            wireless-80211;
          }
        }
      }
      attributes {
        ignore {
          framed-ip-netmask;
        }
        exclude {
          accounting-delay-time [accounting-start accounting-stop];
          accounting-session-id [access-request accounting-on accounting-off
          accounting-start accounting-stop];
          dhcp-gi-address [access-request accounting-start accounting-stop];
          dhcp-mac-address [access-request accounting-start accounting-stop];
          nas-identifier [access-request accounting-start accounting-stop];
          nas-port [accounting-start accounting-stop];
```

Copyright © 2017, Juniper Networks, Inc.

```
            nas-port-id [accounting-start accounting-stop];
            nas-port-type [access-request accounting-start accounting-stop];
          }
        }
      }
    }
[edit logical-systems isp-bos-metro-12 routing-instances isp-cmbrg-12-32]
interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.100/24;
            }
        }
    }
    ge-0/0/0 {
        vlan-tagging;
        unit 0 {
            vlan-id 200;
            family inet {
                unnumbered-address lo0.0;
            }
        }
    }
}
```

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 5

Configuring RADIUS Accounting for Subscriber Access

- RADIUS Accounting Statistics for Subscriber Access Overview on page 97
- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting on page 101
- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 102
- RADIUS Acct-On and Acct-Off Messages on page 105
- Configuring Per-Subscriber Session Accounting on page 105
- Configuring Per-Service Session Accounting on page 108
- Configuring Service Packet Counting on page 109
- Configuring Back-up Options for RADIUS Accounting on page 111
- Configuring Service Interim Accounting on page 112

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.

Copyright © 2017, Juniper Networks, Inc.

The router uses the RADIUS attributes and Juniper Networks VSAs listed in Table 14 on page 98 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 14: RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting

Attribute Number	Attribute Name	Type of Statistics
26–151	IPv6-Acct-Input-Octets	IPv6
26–152	IPv6-Acct-Output-Octets	IPv6
26–153	IPv6-Acct-Input-Packets	IPv6
26–154	IPv6-Acct-Output-Packets	IPv6
26–155	IPv6-Acct-Input-Gigawords	IPv6
26–156	IPv6-Acct-Output-Gigawords	IPv6
47	Acct-Input-Packets	IPv4 and IPv6 aggregation
48	Acct-Output-Packets	IPv4 and IPv6 aggregation
52	Acct-Input-Gigawords	IPv4 and IPv6 aggregation
53	Acct-Output-Gigawords	IPv4 and IPv6 aggregation

Related Documentation

- Configuring Per-Subscriber Session Accounting on page 105
- Configuring Per-Service Session Accounting on page 108
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Copyright © 2017, Juniper Networks, Inc.

Understanding RADIUS Accounting Duplicate Reporting

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 15 on page 99 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 15: 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

Copyright © 2017, Juniper Networks, Inc.

Table 15: Duplicate RADIUS Accounting Reporting *(continued)*

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
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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Per-Subscriber Session Accounting on page 105
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting on page 101

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
     ```

Copyright © 2017, Juniper Networks, Inc.

- Specify how the router uses the **exclude** statement configuration to filter RADIUS attributes from accounting interim messages.

  The router uses the configuration for the exclude statement in the duplication access profile to determine which RADIUS attributes are not included in the accounting interim messages.

  ```
  [edit access profile profile-name accounting]
  user@host# set duplication-filter exclude-attributes
  ```

Related Documentation

- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Per-Subscriber Session Accounting on page 105

Preservation of RADIUS Accounting Information During an Accounting Server Outage

If the router loses contact with the RADIUS accounting server, as represented in Figure 1 on page 102, 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 1: Topology with Loss of Access to Accounting Server

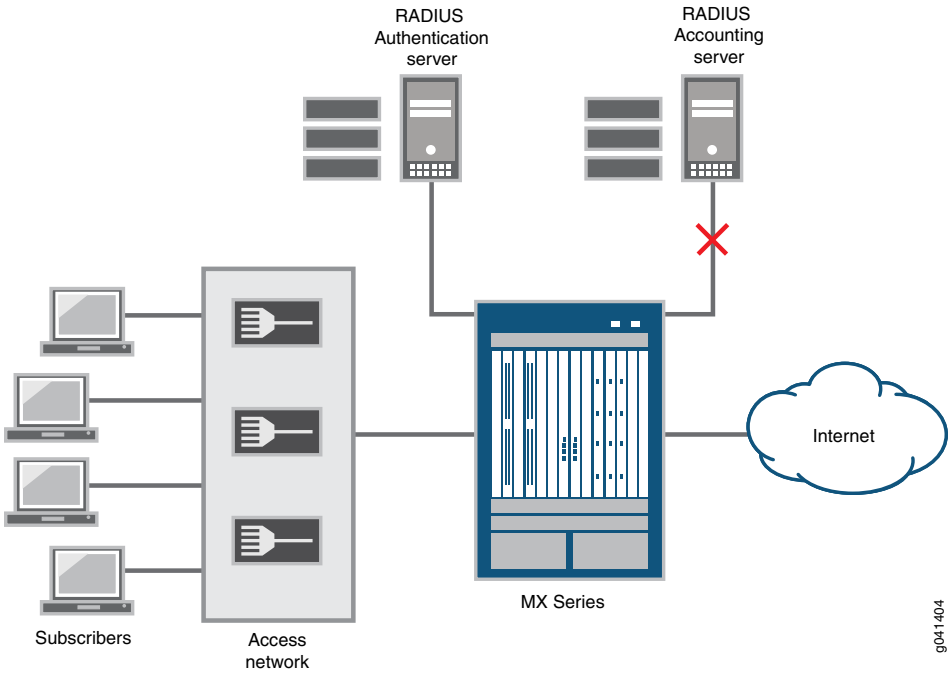

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

The router can hold the pending accounting messages for up to 24 hours. When the configurable maximum holding period passes, all accounting stop messages still in the pending queue are flushed, even if the accounting server has come back online. A consequence of this is that subscriber logins resume immediately if they were failing because the maximum pending limit had been reached.

All pending messages are also flushed in either of the following circumstances:

- If you remove the last accounting server from the access profile, because then there is no place to send the messages.
- If you remove the accounting backup configuration.

While the router is withholding accounting stop messages, you can force the router to attempt contact with the accounting server immediately, rather than allowing it to wait until the periodic interval has expired. When you do so, the router first replays a batch of stop messages to the server, with one of the following outcomes:

- If the router receives an acknowledgment of receipt, then it marks the server as online and begins replaying all remaining pending stop messages in batches.
- If the router does not receive the acknowledgment, then it resumes sending a single pending accounting stop message at the periodic interval.

When a subscriber logs out while the accounting server is offline, the accounting stop requests for the subscriber and the session are queued and replayed to the server when it comes online. In this case, the subscriber session and service session information is retained, so that the router can send a correct accounting request when the server comes back online.

In the event of a graceful Routing Engine switchover while the accounting server is offline, the pending stop messages can be replayed from the active Routing Engine when the server is online again.

NOTE: When RADIUS accounting backup is configured, you must use different servers for RADIUS authentication and accounting. Subscriber authentication fails when the same server is configured for both authentication and accounting.

If the RADIUS server acts on behalf of other back-end RADIUS accounting or authentication servers and forwards requests to them, subscribers can be authenticated but accounting requests are not sent out.

Use the show network-access aaa statistics command to view backup accounting statistics.

Related Documentation

- Configuring Back-up Options for RADIUS Accounting on page 111
- Forcing the Router to Contact the Accounting Server Immediately on page 115

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- AAA Service Framework Overview on page 3
- Configuring Per-Subscriber Session Accounting on page 105
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 39

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 97 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
   ```

Copyright © 2017, Juniper Networks, Inc.

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

Copyright © 2017, Juniper Networks, Inc.

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 an Acct-On message when the first RADIUS server is added to the access profile, and to send 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
    ```

Related Documentation

- RADIUS Accounting Statistics for Subscriber Access Overview on page 97

Copyright © 2017, Juniper Networks, Inc.

- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Per-Service Session Accounting on page 108
- Retaining Authentication and Accounting Information During Session Startup on page 91
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting on page 101
- Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications on page 692
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

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 97 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 16 on page 109 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" on page 109.

Copyright © 2017, Juniper Networks, Inc.

Table 16: 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).

Related Documentation

- Configuring Service Packet Counting on page 109
- RADIUS Accounting Statistics for Subscriber Access Overview on page 97
- Configuring Per-Subscriber Session Accounting on page 105
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Configuring Service Packet Counting

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- *Classic Filters Overview*
- *Defining Dynamic Filter Processing Order*
- RADIUS Accounting Statistics for Subscriber Access Overview on page 97
- Configuring Per-Service Session Accounting on page 108
- Configuring Per-Subscriber Session Accounting on page 105
- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Applying Standard Firewall Filters*
- *Firewall Filter Terminating Actions*
- *Firewall Filter Nonterminating Actions*

Copyright © 2017, Juniper Networks, Inc.

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.

NOTE: Configuring accounting backup disables the revert timer. An error message is generated if you attempt to configure the revert-interval statement at the [edit access profile *profile-name* options] or [edit access radius-options] hierarchy levels.

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]

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set max-pending-accounting-stops 32000
user@host# set max-withhold-time 360
```

Use the show network-access aaa statistics command to view backup accounting statistics.

Related Documentation

- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 102
- Forcing the Router to Contact the Accounting Server Immediately on page 115

Configuring Service Interim Accounting

By default, service accounting is disabled. If you configure service accounting using both RADIUS attributes and the CLI interface, the RADIUS setting takes precedence over the CLI setting.

To configure service accounting for an access profile for a subscriber:

1. At the **[edit access profile *profile-name* service]** hierarchy level, 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 accounting service]
   user@host# set statistics (time | volume-time)
   ```

Related Documentation

- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 6

Configuring Routers and RADIUS Servers for Subscriber Access

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Forcing the Router to Contact the Accounting Server Immediately on page 115
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Options for Subscriber Access Globally on page 120
- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable on page 121

Configuring Router or Switch Interaction with RADIUS Servers

You specify the RADIUS servers that the router or switch can use and you also configure how the router or switch interacts with the specified servers. You can configure the router or switch to connect to multiple RADIUS servers on IPv4 and IPv6 networks. You can configure the RADIUS server connectivity at the **[edit access]** hierarchy level or for a specific dynamic profile at the **[edit access profile]** hierarchy level.

To configure connectivity to a RADIUS server and to specify how the router or switch interacts with the server:

1. Specify the IP address of the RADIUS server.

   The example in this step shows the configuration statement for two RADIUS servers, one with an IPv4 address and the other with an IPv6 address. All other examples in this topic show an IPv4 address. The configuration is similar for IPv6.

   For an IPv4 RADIUS server:

   ```
   [edit access]
   user@host# edit radius-server 192.168.1.250
   ```

   For an IPv6 RADIUS server:

   ```
   user@host# edit radius-server 2001:DB8:0:f101::2
   ```

2. (Optional) Configure the RADIUS server accounting port number. The default accounting port number is 1813.

Copyright © 2017, Juniper Networks, Inc.

```
[edit access radius-server 192.168.1.250]
user@host# set accounting-port 1813
```

3. (Optional) Configure the port number the router or switch uses to contact the RADIUS server. The default port number is 1812.

```
[edit access radius-server 192.168.1.250]
user@host# set port 18914
```

4. Configure the required secret (password) that the local router or switch passes to the RADIUS client. Secrets enclosed in quotation marks can contain spaces.

```
[edit access radius-server 192.168.1.250]
user@host# set secret $ABC123$ABC123
```

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. You can limit the number from 0 through 2000 outstanding requests per RADIUS server. The default setting is 1000 outstanding requests per server.

```
[edit access radius-server 192.168.1.250]
user@host# set max-outstanding-requests 500
```

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 or switch interfaces.

   For an IPv4 source address:

```
[edit access radius-server 192.168.1.250]
user@host# set source-address 192.168.1.100
```

   For an IPv6 source address:

```
[edit access radius-server 2001:DB8:0:f101::2]
user@host# set source-address 2001:DB8:0:f101::1
```

7. (Optional) Configure retry and timeout values for authentication and accounting messages.

   a. Configure how many times the router or switch attempts to contact a RADIUS server when it has received no response. You can configure the router or switch to retry from 1 through 100 times. The default setting is 3 retry attempts.

```
[edit access radius-server 192.168.1.250]
user@host# set retry 4
```

   b. Configure how long the router or switch waits to receive a response from a RADIUS server before retrying the contact. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 1000 seconds.

```
[edit access radius-server 192.168.1.250]
user@host# set timeout 20
```


Copyright © 2017, Juniper Networks, Inc.

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 statements 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 of these options. 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. You can configure the router to retry from 0 through 100 times. The default setting is 0, meaning that this option is disabled.

      ```
      [edit access radius-server 192.168.1.250]
      user@host# set accounting-retry 6
      ```

   b. Configure how long the router waits to receive a response from a RADIUS accounting server before retrying the request. You can configure the timeout to be from 0 through 1000 seconds. The default setting is 0, meaning that this option is disabled.

      ```
      [edit access radius-server 192.168.1.250]
      user@host# set accounting-timeout 20
      ```

Related Documentation

- AAA Service Framework Overview on page 3
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

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

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Related Documentation

- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 102

Configuring RADIUS Server Parameters for Subscriber Access

Include the **radius** statement at the **[edit access profile *profile-name*]** hierarchy level to specify the RADIUS parameters for the subscriber access manager feature. The following list provides an overview of the parameters you can configure:

- The IP addresses of one or more RADIUS authentication and accounting servers.
- Options for the RADIUS servers, such as the following:
  - Format (decimal or description) used for the accounting session
  - Method (round-robin or direct) the router or switch uses to communicate with the servers
  - NAS identifier to use for RADIUS requests
  - Revert time setting that specifies when the router or switch reverts to using the primary RADIUS server
  - Delimiter character and format for the NAS-Port-ID (RADIUS attribute 87) and Calling-Station-ID (RADIUS attribute 31)
- The RADIUS attributes to be ignored or excluded from RADIUS messages.

To configure RADIUS server parameters:

1. Specify that you want to configure RADIUS support.

   ```
   [edit access profile isp-bos-metro-fiber-basic]
   user@host# edit radius
   ```

2. Specify the addresses of RADIUS authentication and accounting servers.

   See "Specifying RADIUS Authentication and Accounting Servers for Subscriber Access" on page 93.

Copyright © 2017, Juniper Networks, Inc.

3. Configure the RADIUS server options.

   See "Configuring RADIUS Server Options for Subscriber Access" on page 117.

4. Configure RADIUS attributes that are ignored or excluded from RADIUS messages.

   See "Configuring How RADIUS Attributes Are Used for Subscriber Access" on page 43.

Related Documentation

- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 93
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43

Configuring RADIUS Server Options for Subscriber Access

You can specify options that the router or switch uses when communicating with RADIUS authentication and accounting servers for subscriber access.

To configure RADIUS authentication and accounting server options:

1. Specify that you want to configure RADIUS.

   ```
   [edit access profile isp-bos-metro-fiber-basic]
   user@host# edit radius
   ```

2. Specify that you want to configure RADIUS options.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius]
   user@host# edit options
   ```

3. (Optional) Configure the method the router or switch uses to access RADIUS accounting servers.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius options]
   user@host# set client-accounting-algorithm round-robin
   ```

4. (Optional) Configure the method the router or switch uses to access RADIUS authentication servers.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius options]
   user@host# set client-authentication-algorithm round-robin
   ```

5. (Optional) Configure the format the router or switch uses to identify the accounting session.

   ```
   [edit access profile isp-bos-metro-fiber-basic radius options]
   user@host# set accounting-session-id-format decimal
   ```

6. (Optional) Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

Copyright © 2017, Juniper Networks, Inc.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set access-loop-id-local
```

7. (Optional) Specify the information that is excluded from the interface description that the router or switch passes to RADIUS for inclusion in RADIUS attribute 87 (NAS-Port-Id).

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set interface-description-format exclude-adapter
```

8. (Optional) Configure the value for the client RADIUS attribute 32 (NAS-Identifier), which is used for authentication and accounting requests.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set nas-identifier 56
```

9. (Optional) Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width of the fields in the NAS-Port attribute. The total of the widths must not exceed 32 bits, or the configuration fails.

   - For Ethernet subscribers:

     ```
     [edit access profile isp-bos-metro-fiber-basic radius options]
     user@host# set nas-port-extended-format ae-width 10 slot-width 4 adapter-width
       2 port-width 4 pw-width 12 stacked-vlan-width 10 vlan-width 2
     ```

     The width value appears in the Cisco NAS-Port-Info AVP (100).

   - For ATM subscribers:

     ```
     [edit access profile retailer01 radius options]
     user@host# set nas-port-extended-format atm slot-width 3 adapter-width 2
       port-width 3 vpi-width 8 vci-width 16
     ```

10. (Optional) Configure the delimiter character that the router inserts between values in RADIUS attribute 87 (NAS-Port-Id).

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set nas-port-id-delimiter %
```

11. (Optional) Configure the information that the router includes in RADIUS attribute 87 (NAS-Port-Id).

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set nas-port-id-format agent-circuit-id agent-remote-id
```

12. (Optional) Configure the delimiter character that the router inserts between values in RADIUS attribute 31 (Calling-Station-Id).

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set calling-station-id-delimiter "%"
```

13. (Optional) Configure the information that the router includes in RADIUS attribute 31 (Calling-Station-Id).

Copyright © 2017, Juniper Networks, Inc.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set calling-station-id-format agent-circuit-id agent-remote-id
```

14. (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 isp-bos-metro-fiber-basic radius options]
user@host# set nas-port-type ethernet wireless-ieee80211
```

NOTE: This statement is ignored if you configure the ethernet-port-type-virtual in the same access profile.

15. (Optional) Configure the router or switch to use a port type of **virtual** to authenticate clients.

```
[edit access profile isp-bos-metro-fiber-basic 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.

16. (Optional) Configure the number of seconds that the router or switch waits after a server has become unreachable.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set revert-interval 259200
```

17. (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 isp-bos-metro-fiber-basic radius options]
user@host# set vlan-nas-port-stacked-format
```

18. (Optional) Configure the router to use the optional behavior when processing CoA requests that include changes to client profile dynamic variables.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set coa-dynamic-variable-validation
```

19. (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 isp-bos-metro-fiber-basic radius options]
```

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set chap-challenge-in-request-authenticator
```

20. (Optional) Configure whether subscribers can successfully log in even when service activation failures related to configuration errors occur during family activation request processing by authd for a newly authenticated subscriber.

```
[edit access profile isp-bos-metro-fiber-basic radius options]
user@host# set service-activation dynamic-profile optional-at-login
```

Related Documentation

- Attaching Access Profiles on page 124
- RADIUS Server Options for Subscriber Access on page 4
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Manual Configuration of the NAS-Port-Type RADIUS Attribute on page 75
- Configuring a NAS-Port-ID with Additional Options on page 63
- Configuring a Calling-Station-ID with Additional Attributes on page 66
- Subscriber Session Logins and Service Activation Failures Overview on page 125
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 94

Configuring RADIUS Options for Subscriber Access Globally

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 number of requests per second that the router can send to all the RADIUS servers collectively.

```
[edit access radius-options]
user@host# set request-rate rate
```

3. (Optional) Configure the number of seconds that the router or switch waits after a server has become unreachable.

```
[edit access radius-options]
user@host# set revert-interval interval
```

4. (Optional) Configure the number of seconds during which a failing authentication server is not marked as down or unreachable.

```
[edit access radius-options]
user@host# set timeout-grace seconds
```


Copyright © 2017, Juniper Networks, Inc.

5. (Optional) Configure a NAS-Port value that is unique across all MX series routers in the network.

```
[edit access radius-options]
user@host# set unique-nas-port chassis-id-width chassis-id-width
user@host# set unique-nas-port chassis-id chassid-id
```

Related Documentation

- Global RADIUS Options for Subscriber Access on page 8

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
```

Related Documentation

- Configuring RADIUS Options for Subscriber Access Globally on page 120
- Global RADIUS Options for Subscriber Access on page 8

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 7

Configuring Access Profiles for Subscriber Access

- Configuring an Access Profile for Subscriber Management on page 123
- Attaching Access Profiles on page 124

Configuring an Access Profile for Subscriber Management

Access profiles enable you to specify subscriber access authentication and accounting parameters.

To configure an access profile:

1. Specify an existing or new access profile name.

   ```
   [edit access]
   user@host# edit profile profile-name
   ```

2. Specify any desired subscriber access authentication and accounting parameters for the access profile.

3. After you create an access profile, you must attach it before it can take effect.

   See "Attaching Access Profiles" on page 124.

Related Documentation

- Attaching Access Profiles on page 124
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 93
- *Configuring Dynamic Authentication for VLAN Interfaces*

Copyright © 2017, Juniper Networks, Inc.

Attaching Access Profiles

After you have created the access profile that specifies the subscriber authentication and accounting parameters, you must specify where the profile is used. Authentication and accounting will not run unless you specify the profile. Subscriber management supports attaching access profiles at the following hierarchy levels:

- **[edit]**
- **[edit routing-instances *routing-instance-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]**

To attach an access profile:

1. Edit the desired hierarchy level.

   ```
   [edit]
   user@host# edit logical-systems LS1 routing-instances RI1
   ```

2. Specify the name of the access profile that you want to attach. For example:

   ```
   [edit logical-systems logical-system-name routing-instances routing-instance-name]
   user@host# set access-profile vz-bos-metro-fios-basic
   ```

Related Documentation

- Configuring an Access Profile for Subscriber Management on page 123
- AAA Service Framework Overview on page 3

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 8

Configuring the Effect of Service Activation Failures on Subscriber Login

- Subscriber Session Logins and Service Activation Failures Overview on page 125
- Configuring How Service Activation Failures Affect Subscriber Login on page 130

Subscriber Session Logins and Service Activation Failures Overview

When a subscriber attempts to log in and is authenticated by RADIUS, the Access-Accept message may include services in the RADIUS Activate-Service VSA (26-65) to be activated for a particular network family. Depending on the configuration and service type, failure to activate a service can result in denial of the subscriber login.

You can use the **service activation** statement at the **[edit access profile** *profile-name* **radius** options] hierarchy level to configure the behavior subsequent to an activation failure.

Use the following options to configure this behavior separately for two types of services:

- **dynamic-profile**—This service type is configured in the dynamic profile that is applied by the subscriber access profile.
- **extensible-service**—This service type is configured in an Extensible Subscriber Services Manager (ESSM) operation script. These services often configure new interfaces for business subscribers.

Use the following options to specify whether successful activation of these services is required or optional for subscriber login access:

- **required-at-login**—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. This is the default behavior for the **dynamic-profile** service type.
- **optional-at-login**—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. This is the default behavior for the **extensible-service** service type.

Copyright © 2017, Juniper Networks, Inc.

NOTE: Failures associated with the activation of subscriber secure policies (for mirroring traffic to a mediation device) have no effect on access by subscribers subject to the policy.

This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policies.

For the **dynamic-profile** service type, configuration errors include the following:

- Parsing errors of the dynamic profile and its attributes.
- Missing mandatory user variables.
- References to dynamic profiles that do not exist.
- Semantic check failures in the dynamic profile.

For the **extensible-service** service type, configuration errors include the following:

- Parsing errors of the operation script.
- Commit failures.

To activate a service, authd sends an activation request for the appropriate services to the subscriber management infrastructure (SMI). For example, if the request is for the IPv4 family, then it requests activation of only the IPv4 services. In turn, the SMI sends requests to the server daemons associated with the service, such as cosd or filterd. The results returned by the daemons determine whether the service activation is a success or a failure.

- When all server daemons report success, then SMI reports success to authd and the service is activated.
- If any server daemon reports a configuration error and no daemons report a nonconfiguration error, then SMI reports a configuration error to authd. The service is not activated, but depending on the configuration, the network family activation may succeed.
- If any server daemon reports a nonconfiguration error, then SMI reports failure to authd and the service is not activated.

Service and Network Family Activation Process

When a subscriber logs in, authd has to activate the corresponding address family after the subscriber is authenticated. The client application, such as DHCP or PPP, can request activation of a single network family, IPv4 or IPv6, or it can sequentially request both families to be activated. Successful network family activation is related to the activation of associated services. The following steps describe the process when authd is configured to use RADIUS for authentication:

Copyright © 2017, Juniper Networks, Inc.

1. A subscriber attempts to log in.
   a. The client application requests authentication from authd.
   b. authd sends an Access-Request message to the RADIUS server.
   c. The RADIUS server sends an Access-Accept message to authd that includes the RADIUS Activate-Service VSA (26-65).
   d. authd caches the service activation attributes and sends a grant to the client application.
2. The client application sends the first Network-Family-Activate request, for either the IPv4 or IPv6 address family. This request is sometimes referred to as the client-activate request.
3. authd reviews the cached service activation attributes and sends an activation request for the appropriate services to the subscriber management infrastructure (SMI). For example, if the request is for the IPv4 family, then it requests activation of only the IPv4 services. In turn, the SMI sends requests to the server daemons associated with the service, such as cosd or filterd.
4. What authd does next depends on whether the service activation request fails and whether the service is optional or required.
   - When the service activation fails due to a configuration error and the service is optional:
     a. authd deletes the cached service activation attributes for the service.

        NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

     b. authd sends an ACK in response to the family activation request and the family is activated.
     c. The subscriber login proceeds.
   - When the service activation fails due to a configuration error and the service is required:
     a. authd deletes the cached service activation attributes for the service.

Copyright © 2017, Juniper Networks, Inc.

NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.

- When the service activation fails due to a nonconfiguration error and the service is either optional or required:

  a. authd deletes the cached service activation attributes for the service.

NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

  b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.

- When the service activation succeeds:

  a. authd activates the service.

  b. authd sends an ACK in response to the family activation request and the family is activated.

  c. The subscriber login proceeds.

5. Unless service activation was required and failed, causing the family activation to fail in the first request, the client application may send a second request, but only for the family not requested the first time. If the first request was for IPv4, then the second request can only be for IPv6. If the first request was for IPv6, then the second request can only be for IPv4.

6. authd reviews the cached service activation attributes and requests activation for the services associated with the requested address family.

7. What authd does next depends on whether the service activation request fails and whether the service is optional or required.

Copyright © 2017, Juniper Networks, Inc.

- When the service activation fails due to a configuration error and the service is optional:

  a. authd deletes the cached service activation attributes for the service.

  NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

  b. authd sends an ACK in response to the family activation request and the family is activated.

  c. The subscriber login proceeds.

- When the service activation fails due to a configuration error and the service is required:

  a. authd deletes the cached service activation attributes for the service.

  NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

  b. authd sends a NACK in response to the family activation. Because this is the second family activation request, the result of the first family activation determines what happens next:

     - If the first family activation was successful and that subscriber logged in, failure of the second request does not halt the current subscriber login. This event also does not cause authd to log out the previous (first request) subscriber.
     - If the first family activation was unsuccessful, failure of the second request causes the client application to terminate the current subscriber login.

- When the service activation fails due to a nonconfiguration error and the service is either optional or required:

  a. authd deletes the cached service activation attributes for the service.

  NOTE: This deletion enables you to re-issue the service request by using a RADIUS change of authorization (CoA) request or a CLI command, without interference from the failed service.

  b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.

Copyright © 2017, Juniper Networks, Inc.

c.

- When the service activation succeeds:

  a. authd activates the service.

  b. authd sends an ACK in response to the family activation request and the family is activated.

  c. The subscriber login proceeds.

Related Documentation

- service-activation on page 1370

Configuring How Service Activation Failures Affect Subscriber Login

You can configure how an activation failure of optional services during subscriber login affects the outcome of the login. These optional services are those referenced by the RADIUS Activate-Service VSA (26-65) that appears in the RADIUS Access-Accept message during the subscriber's initial login.

You can configure these two service-activation types to be required or optional.

- **dynamic-profile**—These services are configured in the dynamic profile that is applied by the subscriber access profile to provide subscriber access and services for broadband applications. By default, service activation is required for successful login. A configuration error during service activation prevents the network family from being activated and causes the subscriber login to fail.
- **extensible-service**—These services are applied by operation scripts handled by the Extensible Subscriber Services Manager (ESSM) daemon (essmd) for business subscribers. By default, service activation is optional for successful subscriber login.

NOTE: **The service-activation statement configuration affects only activation failures due to configuration errors in the dynamic profile or the ESSM operation script. Failures due to nonconfiguration errors always result in denial of access for the subscriber and termination of the login attempt.**

To configure the behavior for dynamic profile services, do one of the following:

- Specify that service activation is optional.

  [edit access profile *profile-name* radius options service-activation]
  user@host# set dynamic-profile optional-at-login

- Specify that service activation is required (the default).

  [edit access profile *profile-name* radius options service-activation]
  user@host# set dynamic-profile required-at-login

Copyright © 2017, Juniper Networks, Inc.

To configure the behavior for ESSM services, do one of the following:

- Specify that service activation is required.

  [edit access profile *profile-name* radius options service-activation]
  user@host# **set extensible-service required-at-login**

- Specify that service activation is optional (the default).

  [edit access profile *profile-name* radius options service-activation]
  user@host# **set extensible-service optional-at-login**

Related Documentation

- Subscriber Session Logins and Service Activation Failures Overview on page 125

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 9

Configuring Session Options for Subscriber Access

- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138
- Removing Inactive Dynamic Subscriber VLANs on page 139

Understanding Session Options for Subscriber Access

You can 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 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

Copyright © 2017, Juniper Networks, Inc.

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.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

- 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.

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 (master) logical system is supported. Because the wholesaler differentiates between multiple retailers by placing each in a different LS:RI, the usernames are appropriately modified for each retailer.

You can select up to eight characters as delimiters to mark the boundary between the discarded and retained portions of the original username; there is no default delimiter. The portion of the name to the right of the selected delimiter is discarded along with the delimiter. By configuring multiple delimiters, a given username structure can result in different modified usernames. You can configure the direction in which the original name is parsed to determine which delimiter marks the boundary. By default, the parse direction is from left to right.

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.

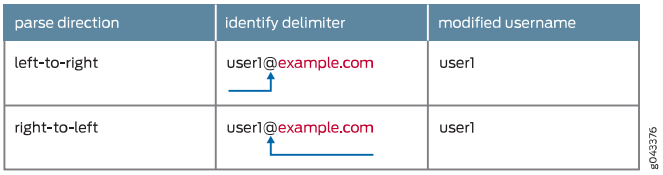

- You specify one delimiter, @. The username is user1@test@example.com. In this case, the parse direction results in different usernames.
  - Parse direction is left-to-right—The left-most @ is identified as the delimiter and test@example.com is discarded. The modified username is user1.
  - Parse direction is right-to-left—The right-most @ is identified as the delimiter and example.com is discarded. The modified username is user1@test.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 14
- Configuring Subscriber Session Timeout Options on page 138
- *Configuring Username Modification for Subscriber Sessions*
- Removing Inactive Dynamic Subscriber VLANs on page 139

Copyright © 2017, Juniper Networks, Inc.

Configuring 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;
      }
    }
  }
}
```

Related Documentation

- Understanding Session Options for Subscriber Access on page 133
- *Configuring Username Modification for Subscriber Sessions*

Copyright © 2017, Juniper Networks, Inc.

- Removing Inactive Dynamic Subscriber VLANs on page 139
- Removing Inactive Dynamic Subscriber VLANs on page 139

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
   ```

Related Documentation

- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138
- client-idle-timeout on page 984

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 10

Receiving DHCP Options From a RADIUS Server

- Centrally Configured Opaque DHCP Options on page 142
- Monitoring DHCP Options Configured on RADIUS Servers on page 146

Copyright © 2017, Juniper Networks, Inc.

Centrally Configured Opaque DHCP Options

Copyright © 2017, Juniper Networks, Inc.

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.

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.

Copyright © 2017, Juniper Networks, Inc.

- Data Flow for RADIUS-Sourced DHCP Options on page 144
- Multiple VSA 26-55 Instances Configuration on page 145
- DHCP Options That Cannot Be Centrally Configured on page 145

Data Flow for RADIUS-Sourced DHCP Options

Figure 2 on page 144 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

Figure 2: DHCP Options Data Flow

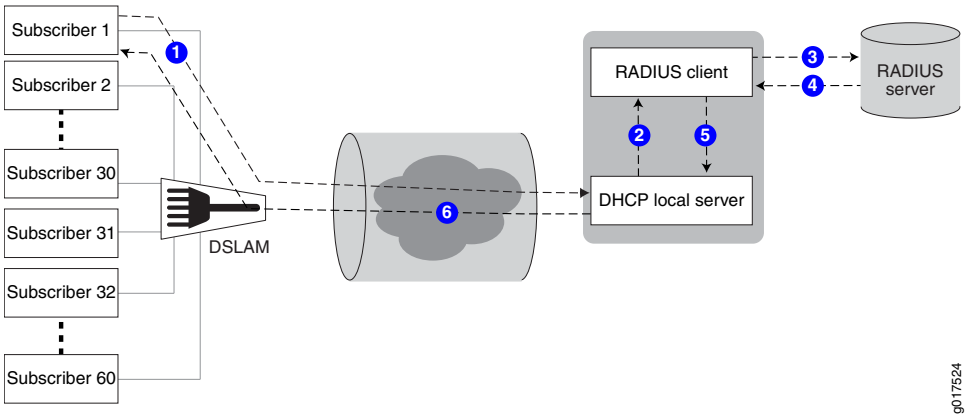

The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

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

Copyright © 2017, Juniper Networks, Inc.

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 17 on page 146 shows the DHCP options that you must not centrally configure on the RADIUS server.

Copyright © 2017, Juniper Networks, Inc.

Table 17: Unsupported Opaque DHCP Options

DHCP Option	Option Name	Comments
Option 0	Pad Option	Not supported.
Option 51	IP Address Lease Time	Value is provided by RADIUS attribute 27 (Session-Timeout).
Option 52	Option Overload	Not supported.
Option 53	DHCP Message Type	Value is provided by DHCP local server.
Option 54	Server Identifier	Value is provided by DHCP local server.
Option 55	Parameter Request List	Value is provided by DHCP local server.
Option 255	End	Value is provided by DHCP local server.
–	DHCP magic cookie	Not supported.

Related Documentation

- Monitoring DHCP Options Configured on RADIUS Servers on page 146

Monitoring DHCP Options Configured on RADIUS Servers

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 detail
Type: 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
```


Copyright © 2017, Juniper Networks, Inc.

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 18 on page 147 describes the first two options in more detail.

Table 18: 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

Related Documentation

- Centrally Configured Opaque DHCP Options on page 142
- show subscribers on page 1820

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 11

Configuring RADIUS Logical Line Identification

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- RADIUS Attributes for LLID Preauthentication Requests on page 150
- Configuring Logical Line Identification (LLID) Preauthentication on page 152
- Configuring a Port and Password for LLID Preauthentication Requests on page 153

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 feature is based on a virtual port — the LLID — rather than the physical line used by the subscriber.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- RADIUS Attributes for LLID Preauthentication Requests on page 150
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

RADIUS Attributes for LLID Preauthentication Requests

Table 19 on page 151 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.

Copyright © 2017, Juniper Networks, Inc.

Table 19: 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.
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**

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

Copyright © 2017, Juniper Networks, Inc.

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 PreAuthLlid
   ```

2. Specify the order in which the router uses the supported preauthentication methods. **radius** is the only supported authentication method.

   ```
   [edit access profile PreAuthLlid]
   user@host# set preauthentication-order radius
   ```

3. Specify that you want to configure RADIUS support.

   ```
   [edit access profile PreAuthLlid]
   user@host# edit radius
   ```

4. Specify the IP address of the RADIUS server used for preauthentication.

   ```
   [edit access profile PreAuthLlid 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
   ```


Copyright © 2017, Juniper Networks, Inc.

```
                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
```

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- RADIUS Attributes for LLID Preauthentication Requests on page 150

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;
}
```

Copyright © 2017, Juniper Networks, Inc.

- To specify the password for a specific access profile:

```
[edit access]
profile profile-name {
   radius-server server-address {
      preauthentication-secret password;
   }
}
```

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- RADIUS Attributes for LLID Preauthentication Requests on page 150

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 12

Provisioning Subscriber Services Using Cisco VSAs

- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155
- Configuring Service Interim Accounting on page 158

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 and it uses Cisco VSAs (Cisco Vendor ID 9) to activate and deactivate the services.

The following Cisco VSAs are used together to activate a service:

- Cisco-AVPair (26-1) with the value of the *subscriber:command=activate-service* parameter
- Cisco-AVPair(26-1) with the value the *subscriber:service-name=service-name* parameter

The following Cisco VSAs are used in together to deactivate a service:

- Cisco-AVPair (26-1) with the value of the *subscriber:command=deactivate-service* parameter
- Cisco-AVPair (26-1) with the value of *subscriber:service-name=service-name* parameter

In such a deployment, you can also enable and disable service accounting and specify service interim accounting interval using the CLI interface because Cisco Broadhop does not send attributes related to service accounting in RADIUS messages to the BNG.

Copyright © 2017, Juniper Networks, Inc.

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.

Table 20 on page 156 describes the service accounting type that becomes effective for different configuration combinations of the service accounting method or type using the local CLI settings and RADIUS attributes:

Table 20: Effective Service Accounting Method with CLI and RADIUS Configurations

Service Accounting-Type Local Configuration	Service Accounting-Type RADIUS Configuration	Service Accounting-Type Applicable
Not configured	Not configured	No service accounting
Not configured	Method configured using RADIUS attributes	Service accounting configured using RADIUS
Method configured using the CLI	Not configured	Service accounting configured using the CLI
Method configured using the CLI	Method configured using RADIUS attributes	Service accounting configured using RADIUS
Method configured using the CLI	Explicitly disabled with a value of 0	No service accounting

Table 21 on page 156 describes the service interim accounting interval that becomes effective for different configuration combinations of the service accounting method or type using the local CLI settings and RADIUS attributes:

Table 21: Effective Service Interim Accounting Interval with CLI and RADIUS Configurations

Service Interim Accounting Interval Using Local Configuration	Service Interim Accounting Interval Using RADIUS Configuration	Service Interim Accounting Interval Applicable
Not configured	Not configured	No service interim accounting
Not configured	Interval configured using RADIUS attributes	Service interim accounting configured using RADIUS
Interval configured using the CLI	Not configured	Service interim accounting configured using the CLI
Interval configured using the CLI	Interval configured using RADIUS attributes	Service interim accounting configured using RADIUS
Interval configured using the CLI	Explicitly disabled with a value of 0	No service interim accounting

Copyright © 2017, Juniper Networks, Inc.

By default, service accounting is disabled. If you configure service accounting using both RADIUS attributes and the CLI interface, the RADIUS setting takes precedence over the CLI setting. To enable service accounting using the CLI, include the **accounting** statement at the **[edit access profile *profile-name* service]** hierarchy level.

```
[edit access]
profile profile-name {
  service {
    accounting;
  }
}
```

To enable interim service accounting updates and configure the amount of time that the router waits before sending a new service accounting update, include the **update-interval *minutes*** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level.

```
[edit access]
profile profile-name {
  service {
    accounting {
      update-interval minutes;
    }
  }
}
```

You can configure the router to collect time statistics, or both volume and time statistics, for the service accounting sessions being managed by AAA. To configure the collection of statistical details that are time-based only, include the **statistics time** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level. To configure the collection of statistical details that are both volume-time-based only, include the **statistics volume-time** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level.

```
[edit access]
profile profile-name {
  service {
    accounting {
      statistics (time | volume-time);
    }
  }
}
```

You can 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 for all of the access profiles:

Copyright © 2017, Juniper Networks, Inc.

```
[edit access]
radius-server server-address {
  dynamic-request-port port-number;
}
```

To specify the UDP port for a specific access profile:

```
[edit access]
profile profile-name {
  radius-server server-address {
    dynamic-request-port port-number;
  }
}
```

Related Documentation

- Configuring Service Interim Accounting on page 112

Configuring Service Interim Accounting

By default, service accounting is disabled. If you configure service accounting using both RADIUS attributes and the CLI interface, the RADIUS setting takes precedence over the CLI setting.

To configure service accounting for an access profile for a subscriber:

1. At the **[edit access profile *profile-name* service]** hierarchy level, 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 accounting service]
   user@host# set statistics (time | volume-time)
   ```

Related Documentation

- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 13

Configuring Domain Maps for Subscriber Management

- Domain Mapping Overview on page 160
- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Configuring a Domain Map on page 164
- Configuring a Wildcard Domain Map on page 165
- Specifying an Access Profile in a Domain Map on page 167
- Specifying an Address Pool in a Domain Map on page 168
- Specifying a Dynamic Profile in a Domain Map on page 168
- Specifying an AAA Logical System/Routing Instance in a Domain Map on page 169
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 170
- Configuring Domain and Realm Name Usage for Domain Maps on page 171
- Specifying Domain and Realm Name Delimiters on page 172
- Specifying the Parsing Order for Domain and Realm Names on page 173
- Specifying the Parsing Direction for Domain and Realm Names on page 173
- Enabling Domain Name Stripping on page 174
- Specifying a Tunnel Profile in a Domain Map on page 175
- Specifying a Tunnel Switch Profile in a Domain Map on page 175
- Configuring PADN Parameters for a Domain Map on page 176

Copyright © 2017, Juniper Networks, Inc.

Domain Mapping Overview

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 an 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.

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 22 on page 161 describes the access options and parameters you can configure in the domain map.

Copyright © 2017, Juniper Networks, Inc.

Table 22: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

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***.

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

- Configuring a Domain Map on page 164

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Domain Mapping Overview on page 160
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 170

Copyright © 2017, Juniper Networks, Inc.

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 165.

   - 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 167.

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 168.

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 168.

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 169.

6. (Optional) Configure the target logical system/routing instance in which AAA requests are sent for the domain map.

Copyright © 2017, Juniper Networks, Inc.

See "Specifying a Target Logical System/Routing Instance in a Domain Map" on page 170.

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 171.

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 174.

9. (Optional) Configure rules to strip the user portion of the username for authentication only. You can also use a directional option to configure a left-to-right (default) or right-to-left application.

   See **strip-username**.

10. (Optional) Specify a password to use for all subscriber authentications. This option affects only the username/password sent in the access-request to external policy/RADIUS servers.

    See **override-password**.

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 175.

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 175.

13. (Optional) Configure the PADN parameters used for PPPoE route information for the domain map.

    See "Configuring PADN Parameters for a Domain Map" on page 176.

Related Documentation

- Domain Mapping Overview on page 160
- Verifying and Managing Domain Map Configuration on page 230

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,

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

See "Configuring a Domain Map" on page 164.

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

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 23 on page 167.

Table 23: 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***

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

Copyright © 2017, Juniper Networks, Inc.

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 24 on page 168.

Table 24: 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***

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

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 25 on page 169.

Copyright © 2017, Juniper Networks, Inc.

Table 25: 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***

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

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***

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Domain Mapping Overview on page 160
- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Configuring a Domain Map on page 164
- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 170

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**)

Copyright © 2017, Juniper Networks, Inc.

NOTE: Subscriber management is supported in the default logical system only.

Related Documentation

- Domain Mapping Overview on page 160
- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Configuring a Domain Map on page 164
- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163

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 173.

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 172.

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 173.

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 174.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

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 \
   ```

Related Documentation
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Configuring Domain and Realm Name Usage for Domain Maps on page 171

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

Copyright © 2017, Juniper Networks, Inc.

produces a domain name of **example.com**. For the same username, left-to-right parsing produces a domain name of **test.com@example.com**.

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
   ```

Related Documentation

- Configuring Domain and Realm Name Usage for Domain Maps on page 171

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
   ```

Related Documentation

- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164
- *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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164
- *Configuring L2TP Tunnel Switching*

Configuring PADN Parameters for a Domain Map

You can configure PPPoE to receive PPPoE Active Discovery Network (PADN) messages when a subscriber connects to a PPPoE server. The PADN information associates the PPPoE session with a set of routes that the session can use. You can configure the route information in a domain map, which enables you to apply specific PADN information to subscribers in a particular domain. You can configure a maximum of 16 routes in a domain map.

To configure PADN parameters 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 PADN route information you want to include in the domain map. For each route, include the destination IP address, subnet mask, and route metric.

```
[edit access domain map domain-map-name]
user@host# set padn destination-address mask destination-mask metric route-metric
```

Related Documentation

- Domain Mapping Overview on page 160
- Configuring a Domain Map on page 164

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 14

Configuring Dynamic Service Activation for Subscriber Access

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- Dynamic Service Activation During Login Overview on page 178
- Configuring RADIUS-Initiated Dynamic Request Support on page 178
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179
- Usage Thresholds for Subscriber Services on page 182
- RADIUS-Initiated Disconnect Overview on page 183
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 184

Using RADIUS Dynamic Requests for Subscriber Access Management

RADIUS dynamic requests provide an efficient way to centrally manage subscriber sessions. The AAA Service Framework's RADIUS dynamic request support allows RADIUS servers to initiate user-related operations, such as a termination operation, by sending unsolicited request messages to the router. Without the RADIUS dynamic request feature, the only way to disconnect a RADIUS user is from the router, which can be cumbersome and time-consuming in large networks.

In a typical client-server RADIUS environment, the router functions as the client and initiates requests sent to the remote RADIUS server. However, when using RADIUS dynamic requests, the roles are reversed. For example, during a disconnect operation, the remote RADIUS server performs as the client and initiates the request (the disconnect action) — the router functions as the server in the relationship.

You create an access profile to configure the router to support RADIUS dynamic requests. This configuration enables the router to receive and act on the following types of messages from remote RADIUS servers:

- Access-Accept messages—Dynamically activate services based on attributes in RADIUS Access-Accept messages received when a subscriber logs in.
- Change-of-Authorization (CoA) messages—Dynamically modify active sessions based on attributes in CoA messages. CoA messages can include service creation requests, deletion requests, RADIUS attributes, and Juniper Networks VSAs.
- Disconnect messages—Immediately terminate specific subscriber sessions.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Dynamic Service Activation During Login Overview on page 178
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179
- RADIUS-Initiated Disconnect Overview on page 183
- Configuring RADIUS-Initiated Dynamic Request Support on page 178
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 13
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 184

Dynamic Service Activation During Login Overview

The AAA Service Framework enables the router to dynamically activate subscriber services as part of a subscriber login operation.

The framework sets up the subscriber session and then completes the service action specified by the Juniper Networks VSA 26–65 that is received in the Access-Accept message. If the service request is unsuccessful, the framework logs out the subscriber.

As part of dynamic service activation, you can also specify usage thresholds for the service—when the threshold is reached, the service is deactivated. You use Juniper Networks VSAs to set the thresholds for the maximum traffic volume for a service and for the length of time the service can be active.

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- Configuring RADIUS-Initiated Dynamic Request Support on page 178
- RADIUS-Initiated Disconnect Overview on page 183
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179
- Usage Thresholds for Subscriber Services on page 182

Configuring RADIUS-Initiated Dynamic Request Support

The router uses the list of specified RADIUS authentication servers for both authentication and dynamic request operations. The router listens on UDP port 3799 for dynamic requests.

To configure RADIUS dynamic request support:

- Specify the IP address of the RADIUS server.

```
[edit access profile isp-bos-metro-fiber-basic radius]
user@host# set authentication-server 192.168.1.3
```

To configure the router to support dynamic requests from more than one RADIUS server:

- Specify the IP addresses of multiple RADIUS servers.

```
[edit access profile isp-bos-metro-fiber-basic radius]
```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# set authentication-server 192.168.1.3 192.168.10.15
```

You can 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 for all of the access profiles:

```
[edit access]
radius-server server-address {
   dynamic-request-port port-number;
}
```

To specify the UDP port for a specific access profile:

```
[edit access]
profile profile-name {
   radius-server server-address {
      dynamic-request-port port-number;
   }
}
```

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- Dynamic Service Activation During Login Overview on page 178
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179
- RADIUS-Initiated Disconnect Overview on page 183
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 13
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 184

RADIUS-Initiated Change of Authorization (CoA) Overview

The AAA Service Framework uses CoA messages to dynamically modify active subscriber sessions. For example, RADIUS attributes in CoA messages might instruct the framework to create, modify, or terminate a subscriber service. You can also use CoA messages to set or modify usage thresholds for current subscriber services.

CoA Messages

Dynamic request support enables the router to receive and process unsolicited CoA messages from external RADIUS servers. RADIUS-initiated CoA messages use the following codes in request and response messages:

- CoA-Request (43)

Copyright © 2017, Juniper Networks, Inc.

- CoA-ACK (44)
- CoA-NAK (45)

Qualifications for Change of Authorization

To complete the change of authorization for a user, you specify identification attributes and session attributes. The identification attributes identify the subscriber. Session attributes specify the operation (activation or deactivation) to perform on the subscriber's session and also include any client attributes for the session (for example, QoS attributes). The AAA Service Framework handles the actual request.

Table 26 on page 180 shows the identification attributes for CoA operations.

NOTE: Using the Acct-Session-ID attribute to identify the subscriber session is more explicit than using the User-Name attribute. When you use the User-Name as the identifier, the CoA operation is applied to the first session that was logged in with the specified username. However, because a subscriber might have multiple sessions associated with the same username, the first session might not be the correct session for the CoA operation.

When you use the Acct-Session-ID attribute, it identifies the specific subscriber session, avoiding that potential error. Although the Acct-Session-ID attribute can include an interface specifier in addition to the session ID—when the attribute is in the description format—only the session ID is used for subscriber matching. For example, if the subscriber has a subscriber session ID of 54785, then the subscriber is matched when the Acct-Session-ID attribute is 54785 (decimal format), or jnpr demux0.1073759682:54785 (description format), or indeed *any value*:54785 (description format).

Table 26: Identification Attributes

Attribute	Description
User-Name [RADIUS attribute 1]	Subscriber username.
Acct-Session-ID [RADIUS attribute 44]	Specific subscriber session.

Table 27 on page 180 shows the session attributes for CoA operations. Any additional client attributes that you include depend on your particular session requirements.

Table 27: Session Attributes

Attribute	Description
Activate-Service [Juniper Networks VSA 26–65]	Service to activate for the subscriber.
Deactivate-Service [Juniper Networks VSA 26–66]	Service to deactivate for the subscriber.

Copyright © 2017, Juniper Networks, Inc.

Table 27: Session Attributes *(continued)*

Attribute	Description
Service-Volume [Juniper Networks VSA 26-67]	Amount of traffic, in MB, that can use the service; service is deactivated when the volume is exceeded.
Service-Timeout [Juniper Networks VSA 26-68]	Number of seconds that the service can be active; service is deactivated when the timeout expires.
Service-Volume-Gigawords [Juniper Networks VSA 26-179]	Amount of traffic, in 4GB units, that can use the service; service is deactivated when the volume is exceeded.
Update-Service [Juniper Networks VSA 26-180]	New values of service and time quotas for existing service.

Message Exchange

The RADIUS server and the AAA Service Framework on the router exchange messages using UDP. The CoA-Request message sent by the RADIUS server has the same format as the Disconnect-Request packet that is sent for a disconnect operation.

The response is either a CoA-ACK or a CoA-NAK message:

- If the AAA Service Framework successfully changes the authorization, the response is a RADIUS-formatted packet with a CoA-ACK message, and the data filter is applied to the session.
- If AAA Service Framework is unsuccessful, the request is malformed, or attributes are missing, the response is a RADIUS-formatted packet with a CoA-NAK message.

NOTE: The AAA Service Framework processes one dynamic request at a time per subscriber. If the framework receives a second dynamic request (either another CoA or a Disconnect-Request) while processing a previous request for the same subscriber, the framework responds with a CoA-NAK message.

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- Dynamic Service Activation During Login Overview on page 178
- RADIUS-Initiated Disconnect Overview on page 183
- Configuring RADIUS-Initiated Dynamic Request Support on page 178
- Usage Thresholds for Subscriber Services on page 182

Copyright © 2017, Juniper Networks, Inc.

Usage Thresholds for Subscriber Services

Starting in Junos OS Release 14.1, subscriber management enables you to manage subscriber services by establishing usage thresholds when a service is dynamically activated or when an existing service is modified by a RADIUS CoA action.The service is deactivated when the specified threshold is reached.

Subscriber management supports two types of usage thresholds—traffic volume and time. You use Juniper Networks VSAs to set the usage thresholds. The VSAs are transmitted in RADIUS Access-Accept messages for dynamically activated services, or in RADIUS-initiated CoA-Request messages for existing services. The volume threshold sets the maximum amount of the total input and output traffic that can use the service before the service is deactivated. A time threshold sets the maximum length of time that the service can be active. Table 28 on page 182 shows the VSAs used for volume and time thresholds.

Table 28: Juniper Network VSAs Used for Service Thresholds

Attribute Number	Attribute Name	Description	Value
26-67	Service-Volume	Amount of input and output traffic, in MB, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8). The router polls the traffic in 10 minute intervals.	• range = 0 through 16777215 MB • 0 = no limit
26-68	Service-Timeout	Number of seconds that the service can be active; service is deactivated when the timeout expires. Tagged VSA, which supports 8 tags (1-8).	• range = 0 through 16777215 seconds • 0 = no timeout
26-179	Service-Volume-Gigawords	Amount of input and output traffic, in 4GB units, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8). The router polls the traffic in 10 minute intervals.	• range = 0 through 16777215 4GB units • 0 = no limit
26-180	Update-Service	New values of service and time quotas for an existing service. Tagged VSA, which supports 8 tags (1-8).	string: *service-name*

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, subscriber management enables you to manage subscriber services by establishing usage thresholds when a service is dynamically activated or when an existing service is modified by a RADIUS CoA action.

Related Documentation

- Dynamic Service Activation During Login Overview on page 178
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179

RADIUS-Initiated Disconnect Overview

This section describes the AAA Service Framework's support for RADIUS-initiated disconnect dynamic requests. The AAA Service Framework uses disconnect messages to dynamically terminate active subscriber sessions.

Disconnect Messages

To centrally control the disconnection of remote access subscribers, the RADIUS dynamic request feature on the router receives and processes unsolicited messages from RADIUS servers.

The dynamic request feature uses the existing format of RADIUS disconnect request and response messages. RADIUS-initiated disconnect uses the following codes in its RADIUS request and response messages:

- Disconnect-Request (40)
- Disconnect-ACK (41)
- Disconnect-NAK (42)

Qualifications for Disconnect

For the AAA Service Framework to disconnect a user, the Disconnect-Request message must contain an attribute with an accounting session ID. The Disconnect-Request message can contain an Acct-Session-Id (44) attribute or an Acct-Multi-Session-Id (50) attribute for the session ID or both. If both the Acct-Session-Id and Acct-Multi-Session-Id attributes are present in the request, the router uses both attributes. If the User-Name (1) attribute is also present in the request, the username and accounting session ID are used to perform the disconnection. The AAA Service Framework handles the actual request.

Message Exchange

The RADIUS server and the AAA Service Framework exchange messages using UDP. The Disconnect-Request message sent by the RADIUS server has the same format as the CoA-Request packet that is sent for a change of authorization operation.

Copyright © 2017, Juniper Networks, Inc.

The disconnect response is either a Disconnect-ACK or a Disconnect-NAK message:

- If the AAA Service Framework successfully disconnects the user, the response is a RADIUS-formatted packet with a Disconnect-ACK message.
- If the AAA Service Framework cannot disconnect the user, the request is malformed, or attributes are missing from the request, the response is a RADIUS-formatted packet with a Disconnect-NAK message.

NOTE: The AAA Service Framework processes one dynamic request at a time per subscriber. If the framework receives a second dynamic request while processing a previous request (either a CoA or another Disconnect-Request) for the same subscriber, the framework responds with a Disconnect-NAK message.

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 177
- Dynamic Service Activation During Login Overview on page 178
- Configuring RADIUS-Initiated Dynamic Request Support on page 178

Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests

When a RADIUS-initiated CoA or disconnect operation is unsuccessful, the router includes an error-cause attribute (RADIUS attribute 101) in the CoA-NAK or Disconnect-NAK message that it sends back to the RADIUS server. If the detected error does not map to one of the supported error-cause attributes, the router sends the message without an error-cause attribute. Table 29 on page 184 describes the error-cause codes.

Table 29: Error-Cause Codes (RADIUS Attribute 101)

Code	Value	Description
401	Unsupported attribute	The request contains an attribute that is not supported (for example, a third-party attribute).
402	Missing attribute	A critical attribute (for example, the session identification attribute) is missing from a request.
404	Invalid request	Some other aspect of the request is invalid, such as if one or more attributes are not formatted properly.
503	Session context not found	The session context identified in the request does not exist on the router.
504	Session context not removable	The subscriber identified by attributes in the request is owned by a component that is not supported.
506	Resources unavailable	A request could not be honored due to lack of available NAS resources (such as memory).

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- RADIUS-Initiated Change of Authorization (CoA) Overview on page 179
- RADIUS-Initiated Disconnect Overview on page 183

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 15

Configuring Termination Causes for Protocols

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189
- AAA Termination Causes and Code Values on page 190
- DHCP Termination Causes and Code Values on page 192
- L2TP Termination Causes and Code Values on page 193
- PPP Termination Causes and Code Values on page 213
- VLAN Termination Causes and Code Values on page 221

Understanding Session Termination Causes and RADIUS Termination Cause Codes

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

Copyright © 2017, Juniper Networks, Inc.

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 30 on page 188 lists the RFC-defined standard RADIUS Acct-Terminate-Cause codes and the corresponding causes.

Table 30: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 30: RFC-Defined Code Values and Termination Causes *(continued)*

Code Value	Termination Cause	Description
15	Service Unavailable	NAS was unable to provide the requested service.
16	Callback	NAS is terminating the current session in order to perform callback for a new session.
17	User Error	Error in the user input caused the session to be terminated.
18	Host Request	Login host terminated the session normally.

Related Documentation

- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189
- AAA Termination Causes and Code Values on page 190
- DHCP Termination Causes and Code Values on page 192
- L2TP Termination Causes and Code Values on page 193
- PPP Termination Causes and Code Values on page 213
- VLAN Termination Causes and Code Values on page 221

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.

Copyright © 2017, Juniper Networks, Inc.

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

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- AAA Termination Causes and Code Values on page 190
- DHCP Termination Causes and Code Values on page 192
- L2TP Termination Causes and Code Values on page 193
- PPP Termination Causes and Code Values on page 213
- VLAN Termination Causes and Code Values on page 221

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

Copyright © 2017, Juniper Networks, Inc.

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 31 on page 191 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 31: Default Mapping Between AAA Termination Causes and Code Values

Internal AAA Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 31: Default Mapping Between AAA Termination Causes and Code Values *(continued)*

Internal AAA Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
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	Admin Reset
shutdown-idle-timeout	4	Idle Timeout
shutdown-remote-reset	10	NAS Request
shutdown-session-timeout	5	Session Timeout

Related Documentation

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

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.

Copyright © 2017, Juniper Networks, Inc.

Table 32 on page 193 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 32: 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

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

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.

Copyright © 2017, Juniper Networks, Inc.

Table 33 on page 194 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 33: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 id	10	NAS Request
session rx iccn avp missing mandatory proxy authen name	10	NAS Request
session rx iccn avp missing mandatory proxy authen response	10	NAS Request
session rx iccn avp missing random vector	10	NAS Request
session rx iccn avp missing secret	10	NAS Request
session rx iccn avp unknown	10	NAS Request
session rx iccn no resources	10	NAS Request
session rx iccn unexpected	10	NAS Request
session rx icrp avp bad hidden	10	NAS Request
session rx icrp avp bad value assigned session id	10	NAS Request
session rx icrp avp duplicate value assigned session id	10	NAS Request
session rx icrp avp malformed bad length	10	NAS Request
session rx icrp avp malformed truncated	10	NAS Request
session rx icrp avp missing mandatory assigned session id	10	NAS Request
session rx icrp avp missing random vector	10	NAS Request
session rx icrp avp missing secret	10	NAS Request
session rx icrp avp unknown	10	NAS Request
session rx icrp no resources	10	NAS Request
session rx icrp unexpected	10	NAS Request

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx occn avp missing mandatory 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
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
session rx ocrq admin close	10	Admin Reset
session rx ocrq authenticate failed host	10	NAS Request

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
session rx wen avp missing secret	10	NAS Request
session rx wen avp unknown	10	NAS Request
session rx wen no resources	10	NAS Request
session timeout connection	10	NAS Request
session timeout inactivity	4	idle timeout
session timeout session	5	session timeout
session timeout upper create	9	NAS Error
session transmit speed unavailable	9	NAS error
session tunnel down	15	Service Unavailable
session tunnel failed	15	Service Unavailable
session tunnel switch profile deleted	6	Admin Reset
session tunneled interface down	8	Port Error
session unknown cause	9	NAS Error
session upper create failed	9	NAS Error
session upper removed	15	Service Unavailable
session warmstart not operational	15	Service Unavailable
session warmstart recovery error	15	Service Unavailable
session warmstart upper not restacked	10	NAS request
tunnel admin close	6	Admin Reset
tunnel admin drain	6	Admin Reset
tunnel control channel failed	15	Service Unavailable

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx scccn avp missing random vector	15	Service Unavailable
tunnel rx scccn avp missing secret	15	Service Unavailable
tunnel rx scccn avp unexpected challenge response	15	Service Unavailable
tunnel rx scccn avp unknown	15	Service Unavailable
tunnel rx scccn no resources	15	Service Unavailable
tunnel rx scccn session id not null	15	Service Unavailable
tunnel rx scccn unexpected	15	Service Unavailable
tunnel rx sccrp authenticate failed challenge	17	User Error
tunnel rx sccrp authenticate failed host	17	User Error
tunnel rx sccrp avp bad hidden	15	Service Unavailable
tunnel rx sccrp avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx sccrp avp bad value bearer capabilities	15	Service Unavailable
tunnel rx sccrp avp bad value challenge	15	Service Unavailable
tunnel rx sccrp avp bad value challenge response	15	Service Unavailable
tunnel rx sccrp avp bad value failover capability	15	Service Unavailable
tunnel rx sccrp avp bad value framing capabilities	15	Service Unavailable
tunnel rx sccrp avp bad value protocol version	15	Service Unavailable
tunnel rx sccrp avp bad value receive window size	15	Service Unavailable
tunnel rx sccrp avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx sccrp avp malformed bad length	15	Service Unavailable
tunnel rx sccrp avp malformed truncated	15	Service Unavailable

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrp avp missing 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx sccrq avp bad value framing capabilities	15	Service Unavailable
tunnel rx sccrq avp bad value protocol version	15	Service Unavailable
tunnel rx sccrq avp bad value receive window size	15	Service Unavailable
tunnel rx sccrq avp duplicate value assigned tunnel id	15	Service Unavailable
tunnel rx sccrq avp malformed bad length	15	Service Unavailable
tunnel rx sccrq avp malformed truncated	15	Service Unavailable
tunnel rx sccrq avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx sccrq avp missing mandatory framing capabilities	15	Service Unavailable
tunnel rx sccrq avp missing mandatory host name	15	Service Unavailable
tunnel rx sccrq avp missing mandatory protocol version	15	Service Unavailable
tunnel rx sccrq avp missing random vector	15	Service Unavailable
tunnel rx sccrq avp missing secret	15	Service Unavailable
tunnel rx sccrq avp unexpected challenge without secret	15	Service Unavailable
tunnel rx sccrq avp unknown	15	Service Unavailable
tunnel rx sccrq bad address	15	Service Unavailable
tunnel rx sccrq no resources	15	Service Unavailable
tunnel rx sccrq no resources max tunnels	15	Service Unavailable
tunnel rx sccrq session id not null	15	Service Unavailable
tunnel rx sccrq unexpected	15	Service Unavailable

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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	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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 malformed bad length	15	Service Unavailable
tunnel rx recovery scccn avp malformed truncated	15	Service Unavailable
tunnel rx recovery scccn avp missing challenge response	17	User Error
tunnel rx recovery scccn avp missing random vector	15	Service Unavailable
tunnel rx recovery scccn avp missing secret	15	Service Unavailable
tunnel rx recovery scccn avp unexpected challenge response	15	Service Unavailable
tunnel rx recovery scccn avp unknown	15	Service Unavailable
tunnel rx recovery scccn no resources	15	Service Unavailable
tunnel rx recovery scccn session id not null	15	Service Unavailable
tunnel rx recovery sccrp authenticate failed challenge	17	User Error
tunnel rx recovery sccrp avp bad hidden	15	Service Unavailable
tunnel rx recovery sccrp avp bad value assigned tunnel id	15	Service Unavailable
tunnel rx recovery sccrp avp bad value bearer capabilities	15	Service Unavailable
tunnel rx recovery sccrp avp bad value challenge	15	Service Unavailable
tunnel rx recovery sccrp avp bad value challenge response	15	Service Unavailable
tunnel rx recovery sccrp avp bad value framing capabilities	15	Service Unavailable
tunnel rx recovery sccrp avp bad value protocol version	15	Service Unavailable

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 33: Default Mapping Between L2TP Termination Causes and Code Values *(continued)*

Internal L2TP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
tunnel rx recovery stopccn avp malformed bad length	15	Service Unavailable
tunnel rx recovery stopccn avp malformed truncated	15	Service Unavailable
tunnel rx recovery stopccn avp missing mandatory assigned tunnel id	15	Service Unavailable
tunnel rx recovery stopccn avp missing mandatory result code	15	Service Unavailable
tunnel rx recovery stopccn avp missing random vector	15	Service Unavailable
tunnel rx recovery stopccn avp missing secret	15	Service Unavailable
tunnel rx recovery stopccn avp unknown	15	Service Unavailable
tunnel rx recovery stopccn no resources	15	Service Unavailable
tunnel rx recovery stopccn session id not null	15	Service Unavailable
tunnel rx recovery unexpected packet	15	Service Unavailable
tunnel rx recovery unknown packet message type indecipherable	15	Service Unavailable
tunnel rx recovery unknown packet message type unrecognized	15	Service Unavailable
tunnel rx session packet null sid invalid	15	Service Unavailable
tunnel rx session packet null sid without assigned session id	15	Service Unavailable
tunnel timeout connection	15	Service Unavailable
tunnel timeout connection recovery tunnel	15	Service Unavailable
tunnel timeout idle	1	User Request
tunnel unknown cause	9	NAS Error

Copyright © 2017, Juniper Networks, Inc.

Table 33: 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 warmstart not operational	15	Service Unavailable
tunnel warmstart recovery error	15	Service Unavailable

Related Documentation

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

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 34 on page 214 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 34 on page 214 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.

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 link selection mismatch	10	NAS Request
bundle fail local mped not set yet	10	NAS Request
bundle fail local mrru mismatch	10	NAS Request
bundle fail local mru mismatch	10	NAS Request
bundle fail peer mrru mismatch	10	NAS Request
bundle fail reassembly location	10	NAS Request
bundle fail reassembly mismatch	10	NAS Request
bundle fail record network	10	NAS Request
bundle fail server location mismatch	10	NAS Request
bundle fail static link	10	NAS Request
failover during authentication	6	Admin Reset
interface admin disable	6	Admin Reset
interface down	2	Lost Carrier
interface no hardware	8	Port Error
ip admin disable	10	NAS Request
ip inhibited by authentication	10	NAS Request
ip link down	10	NAS Request
ip max configure exceeded	10	NAS Request
ip no local ip address	10	NAS Request
ip no local ip address mask	10	NAS Request
ip no local primary dns address	10	NAS Request

Copyright © 2017, Juniper Networks, Inc.

Table 34: Default Mapping Between PPP Termination Causes and Code Values *(continued)*

Internal PPP Termination Cause	RADIUS Acct-Terminate-Cause Attribute	
	Code Value	RADIUS Termination Cause
ip no 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 link down	10	NAS Request
ipv6 local and peer interface ids identical	10	NAS Request
ipv6 max configure exceeded	10	NAS Request
ipv6 no local ipv6 interface id	10	NAS Request
ipv6 no peer ipv6 interface id	10	NAS Request
ipv6 no service	10	NAS Request
ipv6 peer renegotiate rx conf ack	10	NAS Request
ipv6 peer renegotiate rx conf nak	10	NAS Request
ipv6 peer renegotiate rx conf rej	10	NAS Request
ipv6 peer renegotiate rx conf req	10	NAS Request
ipv6 peer terminate code rej	10	NAS Request
ipv6 peer terminate term ack	10	NAS Request
ipv6 peer terminate term req	10	NAS Request
ipv6 service disable	10	NAS Request
ipv6 stale stacking	10	NAS Request
lcp authenticate terminate hold	10	NAS Request
lcp configured mrru too small	10	NAS Request
lcp configured mru invalid	10	NAS Request
lcp configured mru too small	10	NAS Request
lcp dynamic interface hold	10	NAS Request
lcp keepalive failure	10	NAS Request

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 code rej	1	User Request
lcp peer terminate term ack	1	User Request
lcp peer terminate term req	1	User Request
lcp peer terminate protocol reject	1	User Request
lcp peer renegotiate rx conf ack	1	User Request
lcp peer renegotiate rx conf nak	1	User Request
lcp peer renegotiate rx conf rej	1	User Request
lcp peer renegotiate rx conf req	1	User Request
lcp tunnel disconnected	10	NAS Request
lcp tunnel failed	10	NAS Request
link interface no hardware	8	Port Error
lower interface attach failed	2	Lost Carrier
lower interface teardown	2	Lost Carrier
mpls admin disable	10	NAS Request
mpls link down	10	NAS Request
mpls max configure exceeded	10	NAS Request
mpls no service	10	NAS Request
mpls peer renegotiate rx conf ack	10	NAS Request
mpls peer renegotiate rx conf nak	10	NAS Request
mpls peer renegotiate rx conf rej	10	NAS Request
mpls peer renegotiate rx conf req	10	NAS Request

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 34: 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 renegotiate rx conf req	10	NAS Request
osi peer terminate code rej	10	NAS Request
osi peer terminate term ack	10	NAS Request
osi peer terminate term req	10	NAS Request
osi service disable	10	NAS Request
osi stale stacking	10	NAS Request
recovery active state cleanup	9	NAS Error
recovery configured state cleanup	9	NAS Error
recovery init state cleanup	9	NAS Error
recovery terminated state cleanup	9	NAS Error
recovery terminating state cleanup	9	NAS Error
session init failed	9	NAS Error
subscriber mgr activation failed	9	NAS Error
subscriber mgr get credentials failed	9	NAS Error
subscriber mgr link interface not found	9	NAS Error
subscriber mgr set state active failed	9	NAS Error

Related Documentation

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

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

Copyright © 2017, Juniper Networks, Inc.

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 35 on page 222 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 35: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 35: 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 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.
vlan subscriber-activate-error	9	VLAN session termination initiated by an error while attempting to activate the subscriber services for the session.

Related Documentation

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 16

Configuring Extensible Subscriber Services Manager

- Extensible Subscriber Services Manager Overview on page 225
- Understanding the Dictionary File on page 226

Extensible Subscriber Services Manager Overview

Extensible Subscriber Services Manager (ESSM) is a background process that is part of the Intelligent Customer Extendable authentication, authorization, and accounting (ICE-AAA) framework, which supports customer extensible services for both business and residential subscribers. Services are classified as residential or business on the basis of the value specified for the RADIUS VSA (26-173) **ERX-Service-Activate-Type** that is received in the Access-Accept message.

Extensible Subscriber Services Manager uses the ICE-AAA framework, which comprises a dictionary, operation scripts, and RADIUS vendor-specific attributes (VSAs), to create business services for subscribers without modifying Junos OS. Extensible Subscriber Services Manager supports only the **ERX-Activate** service type.

Using the Extensible Subscriber Services Manager, you can create business services using the following sources:

- The dictionary that refers to or invokes the operation scripts.
- The operation scripts that you use to create subscriber-specific configuration
- The VSAs that the RADIUS server sends that contain configuration values for provisioning services

Related Documentation

- Understanding the Dictionary File on page 226
- **show subscribers on page 1820**
- **show subscribers summary on page 1846**
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*

Copyright © 2017, Juniper Networks, Inc.

Understanding the Dictionary File

The XML-based dictionary specifies the action to be taken by ESSMD when it receives a service request. The dictionary contains provisioning, deprovisioning, and operation scripts. ESSMD parses the dictionary file during initialization and stores the parsed information in the database. Extensible Subscriber Services Manager acts on the extensible-subscriber-service requests on the basis of the services configured in the dictionary file.

During a commit operation, essmd verifies the path and the filename of the dictionary file. If the path or the filename is invalid, the commit operation fails and the error is logged in a system log message. Restarting the daemon or performing a graceful Routing Engine switchover (GRES) operation forces essmd to use the new dictionary. Ensure that you always configure a valid dictionary for essmd.

When loading the dictionary file after a successful commit operation, essmd validates whether:

- There are errors in parsing the dictionary file.
- The operation scripts specified in the dictionary file are available on the router.
- Any active services are modified.

If the validation fails, an error is logged in a system log message, and essmd continues to use the existing version of the dictionary file. Use the request services extensible-subscriber-services reload-dictionary command to reload the dictionary file after resolving the errors.

Related Documentation

- dictionary on page 1043
- show extensible-subscriber-services dictionary on page 1712
- show extensible-subscriber-services dictionary attributes on page 1716
- show extensible-subscriber-services dictionary services on page 1719

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 17

Monitoring and Managing AAA Information for Subscriber Access

- Verifying and Managing Subscriber AAA Information on page 227
- Monitoring Pending RADIUS Accounting Stop Messages on page 228
- Verifying and Managing the RADIUS Dynamic-Request Feature on page 229
- Verifying and Managing Domain Map Configuration on page 230
- Verifying and Managing LLID Preauthentication Configuration on page 230

Verifying and Managing Subscriber AAA Information

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 AAA accounting statistics:

  user@host> clear network-access aaa statistics **accounting**

- To clear AAA address-assignment statistics for a client:

Copyright © 2017, Juniper Networks, Inc.

```
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

- CLI Explorer

Monitoring Pending RADIUS Accounting Stop Messages

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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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
```

Related Documentation

- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 102
- Configuring Back-up Options for RADIUS Accounting on page 111

Verifying and Managing the RADIUS Dynamic-Request Feature

Purpose Display RADIUS dynamic request statistics and information.

Action

- To display RADIUS dynamic request statistics:

Copyright © 2017, Juniper Networks, Inc.

user@host>show network-access aaa statistics **dynamic-requests**

Related Documentation
- CLI Explorer

Verifying and Managing Domain Map Configuration

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 on page 160
- Configuring a Domain Map on page 164

Verifying and Managing LLID Preauthentication Configuration

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

Related Documentation
- RADIUS Logical Line Identifier (LLID) Overview on page 149
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

Copyright © 2017, Juniper Networks, Inc.

PART 2

Configuring DHCP for Subscriber Management

- Using DHCP Overview on page 233
- Configuring Default Services That are Activated at Subscriber Login on page 247
- Assigning IP Addresses on page 249
- Configuring Lease Times for IP Addresses on page 257
- Configuring DHCP Individual and Bulk Leasequery on page 265
- Requesting DHCP Client Configuration Information From an Address Pool on page 277
- Authenticating DHCP Clients Using An External AAA Authentication Service on page 281
- Grouping Interfaces and Applying a Common DHCP Configuration to the Group on page 287
- Configuring the Number of DHCP Clients Per Interface on page 297
- Maintaining Subscribers During Interface Delete Events on page 301
- Forcing Dynamic Reconfiguration of Clients From a DHCP Local Server on page 305
- Conserving IP Addresses Using DHCP Auto Logout on page 313
- Overriding Default DHCP Local Server Configuration Settings on page 321
- Configuring DHCP Renegotiation While in Bound State on page 329
- Attaching Access Profiles to DHCP Interfaces on page 331
- Applying a Common DHCP Relay Agent Configuration to Groups of DHCP Servers on page 335
- Changing the Gateway IP Address (giaddr) Field and DHCP Relay Request and Release Packet Source Address on page 337
- Securing DHCP Messages Sent Between DHCP Clients and Servers in Different VRFs on page 339
- Configuring DHCP Relay Agent on page 345
- Configuring DHCP Relay Proxy Mode on page 369
- Configuring DHCP Local Server Authentication on page 371
- Configuring a Minimum DHCP Local Server Configuration on page 373
- Protecting the Routing Engine Using DHCP Firewall Filters on page 375
- Monitoring and Managing DHCP on page 381

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 18

Using DHCP Overview

- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240
- DHCP Relay Proxy Overview on page 243
- DHCP Attributes Overview on page 244

Copyright © 2017, Juniper Networks, Inc.

Extended DHCP Local Server Overview

Copyright © 2017, Juniper Networks, Inc.

Junos OS includes an extended DHCP local server that enhances traditional DHCP server operation by providing additional address assignment and client configuration functionality and flexibility in a subscriber-aware environment. The 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. The address-assignment pools are considered external because they are external to the DHCP local server. The pools are managed independently of the DHCP local server, and can be shared by different client applications, such as DHCP or PPPoE access. Table 36 on page 235 provides a comparison of the extended DHCP local server and a traditional DHCP local server.

The extended DHCP local server provides an IP address and other configuration information in response to a client request. The server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide DHCP client authentication. You can configure the dynamic profile and authentication support on a global basis or for a specific group of interfaces.

Table 36: Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server

Feature	Extended DHCP Local Server	Traditional 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

You can also configure the extended DHCP local server to support 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.

Copyright © 2017, Juniper Networks, Inc.

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.

This overview covers:

- Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools on page 236
- Providing DHCP Client Configuration Information on page 237
- Minimal Configuration for Clients on page 238
- DHCP Local Server and Address-Assignment Pools on page 238

Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools

The pattern of interaction between the DHCP local server, the DHCP client, and address-assignment pools is the same regardless of whether you are using a router or a switch. However, there are some differences in the details of usage.

- On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer or customer premises equipment (CPE), and the DHCP local server is configured on the router.
- On switches—In a typical network configuration, the DHCP client is on an access device, such as a personal computer, and the DHCP local server is configured on the switch.

The following steps provide a high-level description of the interaction among the DHCP local server, DHCP client, and address-assignment pools:

1. The DHCP client sends a discover packet to one or more DHCP local servers in the network to obtain configuration parameters and an IP address for the subscriber (or DHCP client).

2. Each DHCP local server that receives the discover packet then searches its address-assignment pool for the client address and configuration options. Each local server creates an entry in its internal client table to keep track of the client state, then sends a DHCP offer packet to the client.

3. On receipt of the offer packet, the DHCP client selects the DHCP local server from which to obtain configuration information and sends a request packet indicating the DHCP local server selected to grant the address and configuration information.

4. The selected DHCP local server sends an acknowledgement packet to the client that contains the client address lease and configuration parameters. The server also installs the host route and ARP entry, and then monitors the lease state.

Copyright © 2017, Juniper Networks, Inc.

Providing DHCP Client Configuration Information

When the extended DHCP application receives a response from an external authentication server, the response might include information in addition to the IP address and subnet mask. The extended DHCP application uses the information from the authentication grant for the response the DHCP application sends to the DHCP client. The DHCP application can either send the information in its original form or the application might merge the information with local configuration specifications. For example, if the authentication grant includes an address pool name and a local configuration specifies DHCP attributes for that pool (such as, DNS server address), the extended DHCP application merges the authentication results and the attributes in the reply that the server sends to the client.

A local configuration is optional — a client can be fully configured by the external authentication service. However, if the external authentication service does not provide client configuration, you might need to configure the local address-assignment pool to provide the configuration information, such as DNS server, for the client. When a local configuration specifies options, the extended DHCP application adds the local configuration options to the offer PDU the server sends to the client. If the two sets of options overlap, the options in the authentication response from the external service take precedence.

When you use RADIUS to provide the authentication, the additional information might be in the form of RADIUS attributes and Juniper Networks VSAs. Table 37 on page 237 lists the information that RADIUS might include in the authentication grant. See "RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework" on page 13 for a complete list of RADIUS attributes and Juniper Networks VSAs that the extended DHCP applications supports for subscriber access management or DHCP management.

Table 37: Information in Authentication Grant

Attribute Number	Attribute Name	Description
RADIUS attribute 8	Framed-IP-Address	Client IP address
RADIUS attribute 9	Framed-IP-Netmask	Subnet mask for client IP address (DHCP option 1)
Juniper Networks VSA 26-4	Primary-DNS	Primary domain server (DHCP option 6)
Juniper Networks VSA 26-5	Secondary-DNS	Secondary domain server (DHCP option 6)
Juniper Networks VSA 26-6	Primary-WINS	Primary WINS server (DHCP option 44)
Juniper Networks VSA 26-7	Secondary-WINS	Secondary WINS server (DHCP option 44)

Copyright © 2017, Juniper Networks, Inc.

Table 37: Information in Authentication Grant *(continued)*

Attribute Number	Attribute Name	Description
RADIUS attribute 27	Session-Timeout	Lease time
RADIUS attribute 88	Framed-Pool	Address assignment pool name
Juniper Networks VSA 26-109	DHCP-Guided-Relay-Server	DHCP relay server

Minimal Configuration for Clients

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.

DHCP Local Server and Address-Assignment Pools

In the traditional DHCP server operation, the client address pool and client configuration information reside on the DHCP server. With the extended DHCP local server, the client address and configuration information reside in external address-assignment pools (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.

The extended DHCP local server also supports advanced pool matching and the use of named address ranges. You can also 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.

Configuring the DHCP environment that includes the extended DHCP local server requires two independent configuration operations, which you can complete in any order. In one operation, you configure the extended DHCP local server on the router and specify how the DHCP local server determines which address-assignment pool to use. In the other operation, you configure the address-assignment pools used by the DHCP local server.

Copyright © 2017, Juniper Networks, Inc.

The address-assignment pools contain the IP addresses, named address ranges, and configuration information for DHCP clients.

NOTE: **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.**

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- *Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview*
- Using External AAA Authentication Services with DHCP on page 281
- Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option on page 253
- *Graceful Routing Engine Switchover for DHCP*
- *High Availability Using Unified ISSU in the PPP Access Network*
- Tracing Extended DHCP Operations on page 829
- Verifying and Managing DHCP Local Server Configuration on page 385
- Example: Minimum Extended DHCP Local Server Configuration on page 373
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 255
- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 375

Copyright © 2017, Juniper Networks, Inc.

Extended DHCP Relay Agent Overview

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 455 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]**

This overview covers:

Copyright © 2017, Juniper Networks, Inc.

- Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers on page 241
- DHCP Liveness Detection on page 242

Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers

The pattern of interaction among the DCHP 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

Copyright © 2017, Juniper Networks, Inc.

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.**

Related Documentation

- DHCPv6 Relay Agent Overview on page 455
- *Access and Access-Internal Routes for Subscriber Management*
- *Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview*
- Using External AAA Authentication Services with DHCP on page 281
- DHCP Relay Proxy Overview on page 243
- *Graceful Routing Engine Switchover for DHCP*
- *High Availability Using Unified ISSU in the PPP Access Network*
- Verifying and Managing DHCP Relay Configuration on page 385
- Tracing Extended DHCP Operations on page 829
- Example: Minimum DHCP Relay Agent Configuration on page 360
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers on page 365
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 361
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 291
- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 375

Copyright © 2017, Juniper Networks, Inc.

DHCP Relay Proxy Overview

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits.

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.

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.

NOTE: **You cannot configure both DHCP relay proxy and extended DHCP local server on the same interface.**

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Enabling DHCP Relay Proxy Mode on page 369
- *Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity*

DHCP Attributes Overview

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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

- Configuring DHCP Attributes Outside Address Assignment Pools on page 327

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 19

Configuring Default Services That are Activated at Subscriber Login

- Default Subscriber Service Overview on page 247
- Configuring a Default Subscriber Service on page 248

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 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring a Default Subscriber Service on page 248
- CLI-Activated Subscriber Services on page 609
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 731

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
     ```

Related Documentation

- Default Subscriber Service Overview on page 247
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 20

Assigning IP Addresses

- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option on page 253
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- Specifying the Subnet for DHCP Client Address Assignment on page 255
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 255
- DNS Address Assignment Precedence on page 256

DHCP Attributes for Address-Assignment Pools

This topic provides descriptions of DHCPv4 and HCPv6 options.

Table 38 on page 249 describes the DHCPv4 client attributes that you can configure.

Table 38: 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

Copyright © 2017, Juniper Networks, Inc.

Table 38: DHCP Attributes *(continued)*

Attribute	Description	DHCP Option
netbios-node-type	NetBIOS node type.	46
option	User-defined options.	–
option-match	Option 82 value is mapped to named address range.	–
router	IP address for routers on the subnetwork.	3
server-identifier	IP address used as the DHCP source address	54
t1-percentage	Percentage of the **preferred-lifetime** that the client (router) waits before sending renew messages to the DHCPv4 server that granted the original lease to extend the client's lease.	58
t1-renewal-time	Time that the client (router) waits before sending renew messages to extend the client's lease. The renew messages are sent to the DHCPv4 server that granted the original lease. This attribute is an alternative to **t1-percentage**.	58
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 39 on page 250 describes the DHCPv6 client attributes that you can configure.

Table 39: 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.	–

Copyright © 2017, Juniper Networks, Inc.

Table 39: DHCPv6 Attributes *(continued)*

Attribute	Description	DHCPv6 Option
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.	–
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.	–

Related Documentation

- DHCP Attributes Overview on page 244
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254

Copyright © 2017, Juniper Networks, Inc.

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. You use the **pool-match-order** statement to specify the match order. If you do not specify the **pool-match-order**, the router (or switch) uses the default **ip-address-first** matching to select the address pool. After DHCP local server determines the address assignment pool to use, the server performs the matching based on the criteria you specified in the pool configuration.

In the default **ip-address-first** matching, 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 there is no giaddr in the request, then the DHCP local server matches the IP address of the receiving interface to the address of the address-assignment pool.

In **external-authority** matching, 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.

For IPv4 address-assignment pools, you can optionally configure the extended DHCP local server to match the DHCP relay agent information option (option 82) in the client DHCP packets to a named range in the address-assignment pool used for the client. Named ranges are subsets within the overall address-assignment pool address range, which you can configure when you create the address-assignment pool.

NOTE: 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.

To configure the matching order the extended DHCP local server uses to determine the address-assignment pool used for a client:

1. Access the **pool-match-order** configuration.

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

2. Specify the pool matching methods in the order in which the router (switch) performs the methods. You can specify the methods in any order. All methods are optional—the router (switch) uses the **ip-address-first** method by default.

   - Configure the router (switch) to use an external addressing authority.

     ```
     [edit system services dhcp-local-server pool-match-order]
     ```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# set external-authority
```

- Configure the router (switch) to use the ip-address-first method.

```
[edit system services dhcp-local-server pool-match-order]
user@host# set ip-address-first
```

- (IPv4 address-assignment pools only) Specify the option 82 matching method.

```
[edit system services dhcp-local-server pool-match-order]
user@host# set option-82
```

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- Extended DHCP Local Server Overview on page 234
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 255

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.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

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

You use the address-assignment pool feature to include application-specific attributes when clients obtain an address. The client application, such as DHCP, uses the attributes to determine how addresses are assigned, and to also provide optional application-specific 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. "DHCP Attributes for Address-Assignment Pools" on page 249 describes the supported attributes you can configure for IPv4 and IPv6 address-assignment pools.

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 599.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Related Documentation

- Extended DHCP Local Server Overview on page 234
- requested-ip-network-match on page 1341

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;
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
}
```

NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Address-Assignment Pools Overview on page 583

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 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.

Related Documentation

- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses on page 59
- Address-Assignment Pools Overview on page 583

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 21

Configuring Lease Times for IP Addresses

- Configuring a DHCP Lease-Time Threshold on page 257
- DHCP Lease-Time Validation Overview on page 259
- DHCP Lease Timers on page 260
- DHCP Asymmetric Leasing Overview on page 261
- Configuring DHCP Asymmetric Leasing on page 263

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]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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 you do not specify a violation action, DHCP binds the client using the third-party lease but marks the binding as lease-time violating.

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
```

Release History Table

Release	Description
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 Lease-Time Validation Overview on page 259
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

Copyright © 2017, Juniper Networks, Inc.

DHCP Lease-Time Validation Overview

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 40 on page 259.

Table 40: 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	–
Violation action of **drop** occurred, or the DHCP packet was not generated	–	warning

Copyright © 2017, Juniper Networks, Inc.

Table 40: Lease-Time Violation Event Logging *(continued)*

Event	syslog	Extended DHCP Traceoptions
Violation action of **override-lease** occurred (DHCP local server only)	–	warning
Lease-time violation	–	warning

Related Documentation

- Configuring a DHCP Lease-Time Threshold on page 257
- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448
- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455

DHCP Lease Timers

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.

Copyright © 2017, Juniper Networks, Inc.

- 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.

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.

Related Documentation

- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- DHCP Attributes for Address-Assignment Pools on page 249

DHCP Asymmetric Leasing Overview

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. Asymmetric leasing has the following benefits:

- 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).

Copyright © 2017, Juniper Networks, Inc.

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.

Release History Table

Release	Description
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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455
- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448
- Configuring DHCP Asymmetric Leasing on page 263

Copyright © 2017, Juniper Networks, Inc.

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 324 and "Overriding Default DHCP Local Server Configuration Settings" on page 321.

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
   ```

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Overriding the Default DHCP Relay Configuration Settings on page 324
- Overriding Default DHCP Local Server Configuration Settings on page 321
- DHCP Asymmetric Leasing Overview on page 261

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 22

Configuring DHCP Individual and Bulk Leasequery

- DHCP Leasequery and Bulk Leasequery on page 265
- Guidelines for Configuring Support for Leasequery and Bulk Leasequery Operations on page 266
- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270
- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273
- Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration on page 274

DHCP Leasequery and Bulk Leasequery

In a subscriber access network, a DHCP server maintains a significant amount of 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, the DHCP relay also maintains the same lease and route information as the DHCP server, and uses that information when providing subscriber services for the clients.

Starting in Junos OS Release 16.1, subscriber management supports the leasequery feature to enable the DHCP relay to quickly and efficiently obtain the current lease information from a DHCP server. DHCP relay typically uses leasequery to restore the relay's lease information database when the information has been lost; such as after the DHCP relay has been rebooted or replaced. Subscriber management supports two types of leasequery operations—the traditional individual leasequery, which updates lease information related to a specific client, and a more powerful bulk leasequery.

Bulk leasequery enables the DHCP relay agent to retrieve bulk lease information for a group of subscribers from a configured DHCP server in a programmed manner that is more resource-efficient than the individual leasequery method. 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 less resources than the UDP connection used for the individual leasequery process. Bulk leasequery also extends the individual leasequery by providing additional query options

Copyright © 2017, Juniper Networks, Inc.

and functionality, as opposed to the individual leasequery, which queries the individual binding for a known target only.

To configure either type of leasequery operation, you enable support on both the DHCP relay agent and the DHCP server, and you specify the query options you want the query operation to use. Once the options are configured, you invoke the appropriate CLI **request** command to initiate the bulk or individual leasequery operation.

For a DHCPv6 bulk leasequery, you can optionally 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 there are no bound subscribers 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.

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, subscriber management supports the leasequery feature to enable the DHCP relay to quickly and efficiently obtain the current lease information from a DHCP server.

Related Documentation

- Guidelines for Configuring Support for Leasequery and Bulk Leasequery Operations on page 266
- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Guidelines for Configuring Support for Leasequery and Bulk Leasequery Operations

When configuring individual or bulk leasequery support, consider the following guidelines:

- The router supports simultaneous configuration of both leasequery and bulk leasequery, and simultaneous dual-stack configuration for both DHCPv4 and DHCPv6. However, for dual stack environments, you must trigger the DHCPv4 and DHCPv6 leasequery (or bulk leasequery) operations separately.
- DHCP relay agent supports leasequery or bulk leasequery over static and dynamic interfaces.
- DHCP local server supports bulk leasequery on static relay-facing interfaces only.
- DHCP local server listens for bulk leasequery requests from the DHCP relay agent on TCP port 547 for DHCPv6, and on TCP port 67 for DHCPv4.
- Bulk leasequery is not supported for DHCP over PPP/PPPoE.
- 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- DHCP Leasequery and Bulk Leasequery on page 265
- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

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 IPv4 address, MAC address, or client ID. You identify DHCPv6 subscribers by IPv6 address or client ID.

You use the following steps to configure and use the individual leasequery operation.

1. 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 349.

     If the network includes integrated routing and bridging (IRB) interfaces, you must also include the **include-l2-interface-for-irb** 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-l2-interface-for-irb
     ```

   - (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*.

     ```
     [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 458.

     If your network includes integrated routing and bridging (IRB) interfaces, you must also include the **include-l2-interface-for-irb** 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-l2-interface-for-irb
     ```

Copyright © 2017, Juniper Networks, Inc.

2. 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.

   - Specify that you want to configure leasequery options for the DHCP relay agent.

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

   - 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 8
     ```

   - 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 3
     ```

3. 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. The following steps describe the configuration for DHCPv4. For DHCPv6, use the procedure at the **[edit system services dhcp-local-server dhcpv6]** hierarchy level.

   - Enable leasequery support for the DHCP local server.

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

   - (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
     ```

4. Initiate the leasequery operation on the DHCP relay agent.

Copyright © 2017, Juniper Networks, Inc.

See "Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 273. DHCP relay agent uses the specified options when sending the leasequery messages to all configured DHCP local servers in the same logical system/routing instance.

5. (Optional) Display and manage statistics and information for the leasequery operation.

   Use the supported **show** and **clear** commands to manage and display information about the leasequery operation for the DHCP relay agent and the DHCP local server. See "Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration" on page 274.

Related Documentation

- DHCP Leasequery and Bulk Leasequery on page 265
- Configuring and Using DHCP Bulk Leasequery on page 270
- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273
- Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration on page 274

Copyright © 2017, Juniper Networks, Inc.

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.

You use the following steps to configure and use bulk leasequery.

1. 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 349.

     If the network includes integrated routing and bridging (IRB) interfaces, you must also include the **include-l2-interface-for-irb** 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-l2-interface-for-irb
     ```

   - (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*.

     ```
     [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 458.

     If your network includes integrated routing and bridging (IRB) interfaces, you must also include the **include-l2-interface-for-irb** 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-l2-interface-for-irb
     ```

2. 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]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# set accept-max-tcp-connections 6
user@host# set request-max-tcp-connections 7
```

3. 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.

   - Specify that you want to configure bulk leasequery options for the DHCP relay agent.

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

   - 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 5
     ```

   - Specify the number of times DHCP relay attempts the TCP connection with the 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 6
     ```

   - (Optional, DHCPv6 only) Specify the optional automatic trigger. The automatic trigger configures DHCPv6 relay agent to automatically initiate bulk leasequery whenever the jdhcp daemon starts (for example, after a reboot, a graceful Routing Engine switchover, or unified ISSU) and there are no bound subscribers in the session database. When the automatic trigger support is configured, you can still use the CLI command to manually trigger bulk leasequery.

     ```
     [edit forwarding-options dhcp-relay bulk-leasequery]
     user@host# set trigger automatic
     ```

4. 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.

   - Enable bulk leasequery support for the DHCP local server.

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

Copyright © 2017, Juniper Networks, Inc.

- (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 5
  ```

- (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 connection.

  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. By limiting the number of empty replies that the DHCP local server sends, you prevent the connection from being consumed by unauthorized or malicious requesters.

  ```
  [edit system services dhcp-local-server allow-bulk-leasequery]
  user@host# set max-empty-replies 4
  ```

- (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
  ```

- (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 90
  ```

5. Initiate the bulk leasequery operation on the DHCP relay agent.

   DHCP relay agent uses the specified options when sending the bulk leasequery messages to all configured DHCP local servers in the same logical system/routing instance.

   - Automatically initiating bulk leasequery—When the automatic trigger feature is configured, DHCP relay agent initiates the bulk leasequery whenever the jdhcp process starts and there are no bound subscribers in the session database.

Copyright © 2017, Juniper Networks, Inc.

- Manually initiating bulk leasequery—Use the appropriate CLI command to manually initiate bulk leasequery. See "Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database" on page 273.

6. Display and manage statistics and information for the bulk leasequery operation.

   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 Leasequery and Bulk Leasequery Configuration" on page 274.

Related Documentation

- DHCP Leasequery and Bulk Leasequery on page 265
- Configuring and Using DHCP Individual Leasequery on page 267
- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273
- Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration on page 274

Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database

You can manually request that DHCP relay agent initiate an individual leasequery or bulk leasequery operation, which requests current lease information from DHCP local servers.

A leasequery updates the DHCP relay agent's lease database with information for an individual client—for DHCPv4, you can specify clients by IP address, MAC address or client ID. For DHCPv6, you can identify clients by IPv6 prefix or client ID.

A bulk leasequery updates the relay agent's lease database for a group of clients—for DHCPv4, you can specify clients by IPv4 address, MAC address or client ID. For DHCPv6, you can identify clients by IPv6 prefix, client ID, Relay-ID option, Relay Agent Remote-ID option, or IPv6 link address.

When you initiate a leasequery or bulk leasequery, you can optionally specify a particular DHCP server (by address), server group, routing instance, or logical system to be queried.

To initiate a DHCPv4 individual leasequery to update binding information:

- For a DHCPv4 client, based on the client's IP address:

  ```
  user@host> request dhcp relay leasequery 192.168.10.34
  ```

- For a DHCPv4 client, based on the client's MAC address:

  ```
  user@host> request dhcp relay leasequery 00:00:5E:00:53:cd
  ```

- For a DHCPv4 client, based on the client identifier option (option 61):

  ```
  user@host> request dhcp relay leasequery 4/1224
  ```

Copyright © 2017, Juniper Networks, Inc.

To initiate a DHCPv6 individual leasequery to update binding information:

- For a DHCPv6 client, based on the client's IPv6 address:

  user@host> request dhcpv6 relay leasequery 2001:db8:1001::1

- For a DHCPv6 client, based on the client ID option (option 1):

  user@host> request dhcpv6 relay leasequery 4/1224

To initiate a DHCPv4 bulk leasequery to update binding information:

- For a DHCPv4 client, based on the client identifier option (option 61):

  user@host> request dhcp relay leasequery 4/1224

- For a DHCPv4 client, based on the client's MAC address:

  user@host> request dhcp relay bulk-leasequery 00:00:5E:00:53:cd

To initiate a DHCPv6 bulk leasequery to update binding information:

- For a DHCPv6 client, based on the client ID option (option 1):

  user@host> request dhcpv6 relay bulk-leasequery 4/1223

- For DHCPv6 clients, based on IPv6 prefix:

  user@host> request dhcpv6 relay bulk-leasequery 2001:db8:1001::1

- For DHCPv6 clients associated with the DHCPv6 relay identified in the specified Relay-ID option:

  user@host> request dhcpv6 relay bulk-leasequery relay-id 4/2121

- For DHCPv6 clients associated with the specified IPv6 link address:

  user@host> request dhcpv6 relay bulk-leasequery link-address 2001:db8:1001:1::1

- For DHCPv6 clients associated with the specified Relay Agent Remote-ID option:

  user@host> request dhcpv6 relay bulk-leasequery remote-id host:ge-1/0/0:100

NOTE: If you use the **bulk-leasequery** command without specifying a remote ID, by default the DHCPv6 relay queries its one relay ID.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Verifying and Managing DHCP Leasequery and Bulk Leasequery Configuration

Purpose View or clear information about DHCP 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.

Copyright © 2017, Juniper Networks, Inc.

Action

- To display leasequery and bulk 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 and bulk 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 and bulk 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 and bulk 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
  ```

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 23

Requesting DHCP Client Configuration Information From an Address Pool

- DHCP Local Server Handling of Client Information Request Messages on page 277
- Enabling Processing of Client Information Requests on page 278

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 information request that indicates what information is desired. 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 has typically been 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Enabling Processing of Client Information Requests on page 278

Enabling Processing of Client Information Requests

By default, DHCP local server and DHCPv6 local server do not respond to 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.

Configure one or more local address pools if you want to use a local pool rather than one provided by AAA. See "Configuring an Address-Assignment Pool Name and Addresses" on page 587. 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 254 for details about how to configure the information sought by clients that send 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]
     ```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 24

Authenticating DHCP Clients Using An External AAA Authentication Service

- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282
- Configuring Passwords for Usernames on page 285

Using External AAA Authentication Services with DHCP

The extended DHCP local server, including DHCPv6 local server, and the extended DHCP relay agent, including DHCPv6 relay agent, support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application can obtain client addresses and DHCP configuration options from the external AAA authentication server.

NOTE: This section uses the term *extended DHCP application* to refer to both the extended DHCP local server and the extended DHCP relay agent.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user logout directive, the extended DHCP application honors the logout and responds as though it were requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP application supports directed logout using the list of configured authentication servers you specify with the authentication-server statement at the **[edit access profile *profile-name*]** hierarchy level.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

To configure DHCP local server and DHCP relay agent authentication support:

1. Specify that you want to configure authentication options.

Copyright © 2017, Juniper Networks, Inc.

- For DHCP local server:

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

- For DHCP relay agent:

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

- For DHCPv6 local server:

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

- For DHCPv6 relay agent:

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

2. (Optional) Configure a password that authenticates the username to the external authentication service.

   See "Configuring Passwords for Usernames" on page 285.

3. (Optional) Configure optional features to create a unique username.

   See "Creating Unique Usernames for DHCP Clients" on page 282.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Local Server Overview on page 448
- DHCPv6 Relay Agent Overview on page 455

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).

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.

Copyright © 2017, Juniper Networks, Inc.

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***.

  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.

- **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.

Copyright © 2017, Juniper Networks, Inc.

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.

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
```

To configure a unique username:

1. Specify that you want to configure authentication.
   - For DHCP local server:

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

   - For DHCPv6 local server:

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

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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


Copyright © 2017, Juniper Networks, Inc.

NOTE: You can also configure authentication and usernames for groups at additional hierarchy levels. See authentication (DHCP Local Server) and authentication (DHCP Relay Agent).

2. Specify that you want to include optional information in the username. (DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support the **username-include** statement.)

```
[edit system services dhcp-local-server authentication]
user@host# set username-include
```

3. (Optional) Specify the optional information you want to include in the username, for example:

```
[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
```

The previous **username-include** configuration produces this unique username:

wallybrown.0090.1a01.1234.enet@example.com

Related Documentation

- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448
- Extended DHCP Relay Agent Overview on page 240
- Using External AAA Authentication Services with DHCP on page 281

Configuring Passwords for Usernames

You can configure an optional password that the extended DHCP application presents to the external AAA authentication service to authenticate the specified username.

To configure a password that authenticates the username:

1. Specify that you want to configure authentication options.
   - For DHCP local server:

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

   - For DHCPv6 local server:

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

   - For DHCP relay agent:

Copyright © 2017, Juniper Networks, Inc.

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

2. Configure the password. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **password** statement.)

```
[edit system services dhcp-local-server authentication]
user@host# set password $ABC123
```

Related Documentation

- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448
- Extended DHCP Relay Agent Overview on page 240
- Using External AAA Authentication Services with DHCP on page 281
- *Special Requirements for Junos OS Plain-Text Passwords*

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 25

Grouping Interfaces and Applying a Common DHCP Configuration to the Group

- Grouping Interfaces with Common DHCP Configurations on page 287
- Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 288
- Configuring Group-Specific DHCP Local Server Options on page 289
- Configuring Group-Specific DHCP Relay Options on page 290
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 291

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

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

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

2. Create the group and assign a name.

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

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the interface ***interface-name*** statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

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

Copyright © 2017, Juniper Networks, Inc.

```
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
```

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Local Server Overview on page 448
- DHCPv6 Relay Agent Overview on page 455
- Configuring Group-Specific DHCP Local Server Options on page 289
- Configuring Group-Specific DHCP Relay Options on page 290
- Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 288

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.

Copyright © 2017, Juniper Networks, Inc.

- 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
  ```

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 287

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.

Copyright © 2017, Juniper Networks, Inc.

- interface—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- overrides—Override the default configuration settings for the extended DHCP local server. For information, see "Overriding Default DHCP Local Server Configuration Settings" on page 321.

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 287

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 to 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 335.
- 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. 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 324.
- 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 356.

Copyright © 2017, Juniper Networks, Inc.

- 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 349.
- 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 247.

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 287

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

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

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

- Requirements on page 291
- Overview on page 292
- Configuration on page 292
- Verification on page 294

Requirements

This example uses the following hardware and software components:

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

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

- Configure DHCPv6 relay agent.

  See "Extended DHCP Relay Agent Overview" on page 240 and "DHCPv6 Relay Agent Overview" on page 455.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 287.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 287.

Copyright © 2017, Juniper Networks, Inc.

Overview

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

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

Configuration

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

- Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings on page 292
- Results on page 293

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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;
```

Copyright © 2017, Juniper Networks, Inc.

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

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

Verification

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

- Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing on page 294

Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing

Purpose Verify the DHCPv6 relay agent selective traffic processing status.

Copyright © 2017, Juniper Networks, Inc.

Action Display statistics for DHCPv6 relay agent.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              10
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_REPL           0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_FORW           0

Packets forwarded:
    Total                       4
    FWD REQUEST                 2
    FWD REPLY                   2
```

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455
- *DHCP Options and Selective Traffic Processing Overview*
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356
- Grouping Interfaces with Common DHCP Configurations on page 287
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 384
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 361

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 26

Configuring the Number of DHCP Clients Per Interface

- Specifying the Maximum Number of DHCP Clients Per Interface on page 297
- Allowing Only One DHCP Client Per Interface on page 298

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]
     ```

Copyright © 2017, Juniper Networks, Inc.

```
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.

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Allowing Only One DHCP Client Per Interface on page 298
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

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.

Copyright © 2017, Juniper Networks, Inc.

- 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 Auto Logout Overview on page 313
- Automatically Logging Out DHCP Clients on page 315

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 27

Maintaining Subscribers During Interface Delete Events

- Subscriber Binding Retention During Interface Delete Events on page 301
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 302

Subscriber Binding Retention During Interface Delete Events

You can configure the router to maintain DHCP subscribers 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

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 302

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
   ```

Related Documentation

- Subscriber Binding Retention During Interface Delete Events on page 301
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 302

Verifying and Managing the DHCP Maintain Subscribers Feature

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:

Copyright © 2017, Juniper Networks, Inc.

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

- Subscriber Binding Retention During Interface Delete Events on page 301
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 28

Forcing Dynamic Reconfiguration of Clients From a DHCP Local Server

- Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients on page 305
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 310
- Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 311
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 312

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients

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,

Copyright © 2017, Juniper Networks, Inc.

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

Copyright © 2017, Juniper Networks, Inc.

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 41 on page 307 lists the actions taken in response to several different events.

Table 41: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 41: Action Taken for Events That Occur During a Reconfiguration *(continued)*

Event	Action
GRES or DHCP restart occurs.	Reconfiguration process is halted.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

Configuring Dynamic Client Reconfiguration of Extended Local Server Clients

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 DHCP 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
   ```


Copyright © 2017, Juniper Networks, Inc.

3. (Optional) Configure an authentication token. The DHCP local server then includes this token inside the authentication option when it sends forcerenew or reconfigure messages. If the service provider has previously configured the DHCP client with this token, then the client can compare that token against the newly received token, and reject the message if the tokens do not match. This functionality corresponds to RFC 3118, *Authentication for DHCP Messages*, section 4.

   a. For all clients:

      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
      ```

   b. For only the DHCP clients serviced by a group of interfaces:

      For DHCPv4:

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

4. For the DHCPv6 server only, you can include the **strict** statement. By default, the server accepts solicit messages from clients that do not support server-initiated reconfiguration. Including this statement causes the server to discard solicit messages from nonsupporting clients; consequently the server does not bind these clients.

   For all DHCPv6 clients:

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

   For only the DHCPv6 clients serviced by a group of interfaces:

   ```
   [edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
   user@host# set strict
   ```

5. (Optional) Configure how the server attempts reconfiguration.

   See "Configuring Dynamic Reconfiguration Attempts for DHCP Clients" on page 310.

6. (Optional) Configure the response to a failed reconfiguration.

   See "Configuring Deletion of the Client When Dynamic Reconfiguration Fails" on page 311.

7. (Optional) Configure the behavior in response to a RADIUS-initiated disconnect.

Copyright © 2017, Juniper Networks, Inc.

See “Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect” on page 312.

8. (Optional) Configure a token for rudimentary server authentication.

   See “Configuring a Token for DHCP Local Server Authentication” on page 371.

9. (Optional) Initiate reconfiguration of some or all client bindings.

   See “Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings” on page 381.

10. (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 450.

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.

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
   ```

 Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

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-abort
  ```

  For DHCPv6:

  ```
  [edit system services dhcp-local-server dhcpv6 reconfigure]
  user@host# set clear-on-abort
  ```

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.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- clear-on-abort on page 976

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- radius-disconnect on page 1299
- trigger on page 1431

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 29

Conserving IP Addresses Using DHCP Auto Logout

- DHCP Auto Logout Overview on page 313
- Automatically Logging Out DHCP Clients on page 315
- How DHCP Relay Agent Uses Option 82 for Auto Logout on page 316
- DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers on page 317
- Automatically Logging Out DHCPv6 Clients on page 318

DHCP Auto Logout Overview

This topic provides an introduction to the DHCP auto logout feature and includes the following sections:

- Auto Logout Overview on page 313
- How DHCP Identifies and Releases Clients on page 314
- Option 60 and Option 82 Requirements on page 315

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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, 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 "How DHCP Relay Agent Uses Option 82 for Auto Logout" on page 316.

Related Documentation

- Automatically Logging Out DHCP Clients on page 315
- How DHCP Relay Agent Uses Option 82 for Auto Logout on page 316
- Allowing Only One DHCP Client Per Interface on page 298
- Clearing DHCP Bindings for Subscriber Access on page 382

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
     ```

Copyright © 2017, Juniper Networks, Inc.

- 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.

Related Documentation

- DHCP Auto Logout Overview on page 313
- How DHCP Relay Agent Uses Option 82 for Auto Logout on page 316
- Allowing Only One DHCP Client Per Interface on page 298
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

How DHCP Relay Agent Uses Option 82 for Auto Logout

Table 42 on page 316 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 right column.

Table 42: DHCP Relay Agent Option 82 Value for Auto Logout

DHCP Relay Agent Configuration Settings					
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"	giaddr in non-snooped packet	Action Taken
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

Copyright © 2017, Juniper Networks, Inc.

Table 42: DHCP Relay Agent Option 82 Value for Auto Logout *(continued)*

DHCP Relay Agent Configuration Settings					
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"	giaddr in non-snooped packet	Action Taken
Yes	Yes	No	–	Zero	Drop packet
Yes	Yes	No	No	Non-zero	Use option 82 from packet
Yes	Yes	No	Yes	Non-zero	Overwrite the configured option 82
Yes	Yes	Yes	No	–	Use option 82 from packet
Yes	Yes	Yes	Yes	–	Overwrite the configured option 82

Related Documentation

- DHCP Auto Logout Overview on page 313
- Automatically Logging Out DHCP Clients on page 315

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Automatically Logging Out DHCPv6 Clients on page 318
- Clearing DHCP Bindings for Subscriber Access on page 382
- Automatically Logging Out DHCP Clients on page 315

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:

Copyright © 2017, Juniper Networks, Inc.

```
[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

- DHCPv6 Match Criteria for Identifying DHCPv6 Subscribers on page 317
- Overriding Default DHCP Local Server Configuration Settings on page 321
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 30

Overriding Default DHCP Local Server Configuration Settings

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Overriding the Default DHCP Relay Configuration Settings on page 324
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327

Overriding 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.

Copyright © 2017, Juniper Networks, Inc.

To override default DHCP local server configuration settings:

1. (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
     ```

     Grouplevel 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
     ```

2. (Optional) Override the maximum number of DHCP clients allowed per interface.

   See "Specifying the Maximum Number of DHCP Clients Per Interface" on page 297.

3. (Optional) Configure DHCP client auto logout.

   See "Automatically Logging Out DHCP Clients" on page 315.

4. (Optional) Enable processing of information requests from clients.

   See "Enabling Processing of Client Information Requests" on page 278.

5. (Optional) Specify that DHCP NAK and FORCERENEW messages support option 82 information.

   See "Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs" on page 340.

6. (Optional, DHCPv6 only) Specify a delegated pool name to use for DHCPv6 multiple address assignment.

Copyright © 2017, Juniper Networks, Inc.

See "Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation" on page 490.

7. (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.

   See "Enabling DHCPv6 Rapid Commit Support" on page 449.

8. (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 603.

9. (Optional, DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

   See "Automatically Logging Out DHCPv6 Clients" on page 318.

10. (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 329.

11. (Optional, DHCPv4 and DHCPv6) Specify that a short lease be sent to the client.

    See "Configuring DHCP Asymmetric Leasing" on page 263.

12. (Optional, DHCPv4 and DHCPv6) Specify DHCP attributes globally or for groups.

    See "Configuring DHCP Attributes Outside Address Assignment Pools" on page 327.

13. (Optional) Delete DHCP override settings.

    See "Deleting DHCP Local Server and DHCP Relay Override Settings" on page 327.

Related Documentation

- Configuring Group-Specific DHCP Local Server Options on page 289
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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 369.

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 337.

4. (DHCPv4 only) Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

   See "Replacing the DHCP Relay Request and Release Packet Source Address" on page 337.

5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.

   See *Overriding Option 82 Information*.

Copyright © 2017, Juniper Networks, Inc.

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 for DHCP Packets" on page 346.

7. (DHCPv4 only) Trust DHCP client packets that have a giaddr of 0 and that contain option 82 information.

   See "Trusting Option 82 Information" on page 346.

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 297.

9. (DHCPv4 only) Configure client auto logout.

   See "DHCP Auto Logout Overview" on page 313.

10. (DHCPv4 and DHCPv6) Enable or disable support for DHCP snooped clients on interfaces.

    See *Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent*.

11. (DHCPv4 and DHCPv6) Delay authentication of subscribers until the DHCP client sends a Request packet.

    See the **delay-authentication**.

12. (DHCPv4 and DHCPv6) Send release messages to the DHCP server when clients are deleted.

    See "Sending Release Messages When Clients Are Deleted" on page 347.

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 329.

14. (DHCPv6 only) Automatically log out existing client when new client solicits on same interface.

    See "Automatically Logging Out DHCPv6 Clients" on page 318.

15. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.

    See "Disabling DHCP Relay" on page 368.

16. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.

Copyright © 2017, Juniper Networks, Inc.

See "Disabling Automatic Binding of Stray DHCP Requests" on page 348.

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 423.

18. (Optional, DHCPv4 and DHCPv6l) Specify that a short lease be sent to the client.

    See "Configuring DHCP Asymmetric Leasing" on page 263.

Related Documentation

- Configuring Group-Specific DHCP Relay Options on page 290
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327

Deleting DHCP Local Server and DHCP Relay Override Settings

You can delete override settings for DHCP local server and DHCP relay globally, for a named group, or for a specific interface within a named group. You can delete a specific override setting or all overrides.

- To delete a specific DHCP override setting at a particular hierarchy level, include the **overrides** statement with the appropriate subordinate statements. For example, to delete the DHCP local server override **interface-client-limit** setting for a group named **marin20**:

  ```
  [edit system services dhcp-local-server]
  user@host# delete group marin20 overrides interface-client-limit
  ```

- To delete all DHCP override settings at a hierarchy level, include the **overrides** statement without any subordinate statements. For example, to delete all DHCP relay overrides for interface **fxp0.0**, which is in group **marin20**:

  ```
  [edit forwarding-options dhcp-relay]
  user@host# delete group marin20 interface fxp0.0 overrides
  ```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

Configuring DHCP Attributes Outside Address Assignment Pools

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. DHCP attributes are often configured to apply only to clients that receive addresses from a specific address assignment pool. You can also configure these attributes to apply to DHCP clients globally or to clients that belong to a group.

Copyright © 2017, Juniper Networks, Inc.

To specify the pool to be used for prefix delegation:

1. Create a DHCP attribute set that you want to apply for 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
```

Related Documentation

- DHCP Attributes Overview on page 244
- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 31

Configuring DHCP Renegotiation While in Bound State

- DHCP Behavior When Renegotiating While in Bound State on page 329

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 receiving a new Discover or Solicit message that has a client ID that matches the existing client.

You can use the **delete-binding-on-renegotiation** statement to configure DHCP local server and relay agent to override the default behavior. 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 tears down the existing entry and sends (DHCP local server) or forwards (relay agent) a new Offer (DHCPv4) or Advertise (DHCPv6) message.

NOTE: In Junos OS releases prior to 15.1, DHCPv6 local server uses the opposite default behavior, and tears down the existing client entry when receiving a Solicit message while in a bound state.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
   ```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 32

Attaching Access Profiles to DHCP Interfaces

- Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview on page 331
- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 332

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.

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.

Related Documentation

- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 332

Copyright © 2017, Juniper Networks, Inc.

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

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 on page 332
- Attaching an Access Profile to a Group of DHCP Subscribers or a Group of DHCP Clients on page 332
- Attaching an Access Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces on page 333

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
    ```


Copyright © 2017, Juniper Networks, Inc.

- 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 287.

To attach an access profile to a group of interfaces:

- At the DHCP configuration hierarchy, specify the name of the interface group and the access profile to attach to the group.
  - For DHCP local server:

    [edit system services dhcp-local-server]
    user@host# **set group boston** interface ***interface-name*** access-profile ***profile-name***

  - For DHCP relay agent:

    [edit forwarding-options dhcp-relay]
    user@host# **set group quebec** interface ***interface-name*** access-profile ***profile-name***

  - For DHCPv6 local server:

    [edit system services dhcp-local-server dhcpv6]
    user@host# **set group boston** interface ***interface-name*** access-profile ***profile-name***

  - For DHCPv6 relay agent:

    [edit forwarding-options dhcp-relay dhcpv6]
    user@host# **set group quebec** interface ***interface-name***access-profile ***profile-name***

Release History Table

Release	Description
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

- Access Profiles for the DHCP Relay Agent and DHCP Local Server Overview on page 331

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 33

Applying a Common DHCP Relay Agent Configuration to Groups of DHCP Servers

- Configuring Named Server Groups on page 335
- Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups on page 335

Configuring Named Server Groups

You can configure a named group of DHCP servers for use by the extended DHCP relay agent on the router or switch.

You specify the name of the DHCP server group and the IP addresses of one or more DHCP servers that belong to this group. You can configure a maximum of five IP addresses per named server group.

To configure a named server group:

1. Specify the name of the server group.

   ```
   [edit forwarding-options dhcp-relay]
   user@host# set server-group myServerGroup
   ```

2. Add the IP addresses of the DHCP servers belonging to the group.

   ```
   [edit forwarding-options dhcp-relay server-group myServerGroup]
   user@host# set 192.168.100.50
   user@host# set 192.168.100.75
   ```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240

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

You can configure an active server group. Using an active server group enables you to apply a common DHCP relay agent configuration to a named group of DHCP server addresses.

Copyright © 2017, Juniper Networks, Inc.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

To configure an active server group:

- Specify the name of the active server group.

  ```
  [edit forwarding-options dhcp-relay]
  user@host# set active-server-group myServerGroup
  ```

To create an active server group as a global DHCP relay agent configuration option, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. To have the group apply only to a named group of interfaces, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level.

Including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level (as a group-specific option) overrides the effect of including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level as a global option.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Grouping Interfaces with Common DHCP Configurations on page 287

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 34

Changing the Gateway IP Address (giaddr) Field and DHCP Relay Request and Release Packet Source Address

- Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent on page 337
- Replacing the DHCP Relay Request and Release Packet Source Address on page 337

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
   ```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

Replacing the DHCP Relay Request and Release Packet Source 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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 35

Securing DHCP Messages Sent Between DHCP Clients and Servers in Different VRFs

- DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs on page 339
- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs on page 340

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 VRFs. The DHCP relay agent can ensure that there is no direct routing between the client virtual routing and forwarding instance (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 Agent Circuit ID (DHCP option 82 suboption 1) for DHCPv4 packets and the Relay Agent Interface-ID (DHCPv6 option 18) for DHCPv6 packets to identify traffic to be relayed.

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

Copyright © 2017, Juniper Networks, Inc.

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.

Release History Table

Release	Description
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 VRFs.

Related Documentation

- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs on page 340
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs

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).

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

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Copyright © 2017, Juniper Networks, Inc.

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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# edit include-option-82 forcerenew nak
   ```

Release History Table

Release	Description
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 Message Exchange Between DHCP Clients and DHCP Server in Different VRFs on page 339
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 36

Configuring DHCP Relay Agent

- DHCP-Initiated Service Change Based on Remote ID on page 345
- Using Layer 2 Unicast Transmission for DHCP Packets on page 346
- Trusting Option 82 Information on page 346
- Sending Release Messages When Clients Are Deleted on page 347
- Disabling Automatic Binding of Stray DHCP Requests on page 348
- Using DHCP Relay Agent Option 82 Information on page 349
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356
- Extracting an Option 82 or Option 37 Substring to Create an Interface Set on page 358
- Configuring DHCP Relay Agent Source Address on page 360
- Example: Minimum DHCP Relay Agent Configuration on page 360
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 361
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers on page 365
- Example: Configuring DHCP-Initiated Service Change Based on Remote ID on page 366
- Disabling DHCP Relay on page 368

DHCP-Initiated Service Change Based on Remote ID

Subscriber management enables you to update a DHCP client's current service through the use of the client's remote ID. The remote ID can be in the option 82, suboption 2 for DHCPv4 clients, or the DHCPv6 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 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.

Copyright © 2017, Juniper Networks, Inc.

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 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.

Related Documentation

- Example: Configuring DHCP-Initiated Service Change Based on Remote ID on page 366

Using Layer 2 Unicast Transmission 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
   ```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

Trusting Option 82 Information

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

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

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]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set send-release-on-delete
```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

Using DHCP Relay Agent Option 82 Information

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.
- Agent Remote ID (suboption 2)—An ASCII string assigned by the DHCP relay agent that securely identifies the client.

Copyright © 2017, Juniper Networks, Inc.

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 456.

The following sections describe the option 82 operations you can configure:

- Configuring Option 82 Information on page 350
- Including a Prefix in DHCP Options on page 352
- Including a Textual Description in DHCP Options on page 354

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:

Copyright © 2017, Juniper Networks, Inc.

(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
     ```

Copyright © 2017, Juniper Networks, Inc.

- 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 352.

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 354.

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 (;).

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

- 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-interface-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

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Copyright © 2017, Juniper Networks, Inc.

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-interface-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
     ```

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:

Copyright © 2017, Juniper Networks, Inc.

```
[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 (configured in the previous step) ASCII string of **video25**, and then forward that traffic to a named local server group:

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

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

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

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

Release History Table

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

Related Documentation

- *DHCP Options and Selective Traffic Processing Overview*
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 361
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 291

Copyright © 2017, Juniper Networks, Inc.

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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
...
```

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

- Configuring Option 82 Information on page 350
- Configuring DHCPv6 Relay Agent Options on page 456
- Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets on page 459
- *Junos OS Predefined Variables*
- *Configuring Predefined Dynamic Variables in Dynamic Profiles*
- *Configuring a Basic Dynamic Profile*

Copyright © 2017, Juniper Networks, Inc.

Configuring DHCP Relay Agent Source Address

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
   ```

Related Documentation

- Configuring Option 82 Information on page 350

Example: 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 {
```


Copyright © 2017, Juniper Networks, Inc.

```
        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.

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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240

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

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

- Requirements on page 361
- Overview on page 361
- Configuration on page 362
- Verification on page 364

Requirements

This example uses the following hardware and software components:

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

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

- Configure DHCP relay agent.

  See "Extended DHCP Relay Agent Overview" on page 240.

- (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 287.

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.

Copyright © 2017, Juniper Networks, Inc.

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

Configuration

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

- Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings on page 362
- Results on page 363

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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;
      }
    }
  }
}
```

Copyright © 2017, Juniper Networks, Inc.

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

Verification

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

- Verifying the Status of DHCP Relay Agent Selective Traffic Processing on page 364

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

Copyright © 2017, Juniper Networks, Inc.

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- *DHCP Options and Selective Traffic Processing Overview*
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 384
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 291

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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
            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

- Extended DHCP Relay Agent Overview on page 240

Example: 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.

Copyright © 2017, Juniper Networks, Inc.

```
[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 Remote ID 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 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
```

Copyright © 2017, Juniper Networks, Inc.

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 Remote ID string, is activated when the client reconnects.

Related Documentation

- DHCP-Initiated Service Change Based on Remote ID on page 345
- Sending Release Messages When Clients Are Deleted on page 347

Disabling DHCP Relay

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

- Extended DHCP Relay Agent Overview on page 240
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 37

Configuring DHCP Relay Proxy Mode

- Enabling DHCP Relay Proxy Mode on page 369

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
   ```

Related Documentation

- DHCP Relay Proxy Overview on page 243
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 38

Configuring DHCP Local Server Authentication

- Configuring a Token for DHCP Local Server Authentication on page 371

Configuring a Token for DHCP Local Server Authentication

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. The client compares the received token with a token already configured on the client. If the tokens do not match, the DHCP client discards the forcerenew message. Use of the token provides rudimentary protection against inadvertently instantiated DHCP servers.

(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 8ysIU9E32k8r
  ```

  For DHCPv6:

  ```
  [edit system services dhcp-local-server dhcpv6 reconfigure]
  user@host# set token 8ysIU9E32k8r
  ```

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.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- token on page 1407

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 39

Configuring a Minimum DHCP Local Server Configuration

- Example: Minimum Extended DHCP Local Server Configuration on page 373

Example: Minimum Extended DHCP Local Server Configuration

This example shows the minimum configuration you need to use for the extended DHCP local server on the router or switch:

```
[edit system services]
dhcp-local-server {
   group group_one {
      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.

This example creates the server group named **group_one**, and specifies that the DHCP local server is enabled on interface **fe-0/0/2.0** within the group. The DHCP local server uses the default pool match configuration of **ip-address-first.**

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.

Related Documentation

- Extended DHCP Local Server Overview on page 234

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 40

Protecting the Routing Engine Using DHCP Firewall Filters

- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 375
- Port Number Requirements for DHCP Firewall Filters on page 379

Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine

This example shows how to configure a firewall filter to ensure that proper DHCP packets can reach the Routing Engine on MX Series routers.

- Requirements on page 375
- Overview on page 375
- Configuration on page 376
- Verification on page 378

Requirements

This configuration example applies only to routers where DHCP local server and DHCP relay agent services are provided by the jdhcpd process rather than the legacy dhcpd process or fud (UDP forwarding) process. MX Series routers, M120 routers, and M320 routers use jdhcpd. For DHCP relay, that means the configuration is required only at the **[edit forwarding-options dhcp-relay]** hierarchy level and not at the **[edit forwarding-options helpers bootp]** hierarchy level.

No special configuration beyond device initialization is required before you can configure this feature.

Overview

Firewall filters that perform some action on DHCP packets at the Routing Engine, such as a filter to protect the Routing Engine by allowing only proper DHCP packets, require that both port 67 (bootps) and port 68 (bootpc) are configured as both source and destination ports.

DHCP packets received on the line cards are encapsulated by jdhcpd with a new UDP header where their source and destination addresses are set to port 68 before being forwarded to the Routing Engine. For DHCP relay and DHCP proxy, packets sent to the

Copyright © 2017, Juniper Networks, Inc.

DHCP server from the router have both the source and destination UDP ports set to 67. The DHCP server responds using the same ports. However, when the line card receives these DHCP response packets, it changes both port numbers from 67 to 68 before passing the packets to the Routing Engine. Consequently the filter needs to accept port 67 for packets relayed from the client to the server, and port 68 for packets relayed from the server to the client.

In this example, you configure two filter terms, **dhcp-client-accept** and **dhcp-server-accept**. The match conditions for **dhcp-client-accept** specify a source address and destination address for broadcast packets, the UDP protocol used for DHCP packets, and the bootpc (68) source port and bootps (67) destination port. Packets that match these conditions are counted and accepted.

The match conditions for **dhcp-server-accept** specify the UDP protocol used for DHCP packets, and both port 67 and 68 for both source port and destination port. Packets that match these conditions are counted and accepted.

NOTE: This example does not show all possible configuration choices, nor does it show how the filter is applied in your configuration. This example applies to both static application of the filter as well as dynamic application with a dynamic profile.

Configuration

CLI Quick Configuration

To quickly configure the sample Routing Engine DHCP filter, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
[edit]
edit firewall family inet filter RE-protect
edit term dhcp-client-accept
set from source-address 0.0.0.0/32
set from destination-address 255.255.255.255/32
set from protocol udp
set from source-port 68
set from destination-port 67
set then count dhcp-client-accept
set then accept
up
edit term dhcp-server-accept
set from protocol udp
set from source-port 67
set from source-port 68
set from destination-port 67
set from destination-port 68
set then count dhcp-server-accept
set then accept
top
```


Copyright © 2017, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a DHCP firewall filter to protect the Routing Engine:

1. Create or specify a firewall filter.

```
[edit firewall]
user@host# edit family inet filter RE-protect
```

2. Create a filter term for the client.

```
[edit firewall family inet filter RE-protect]
user@host# edit term dhcp-client-accept
```

3. Specify the match conditions for DHCP packets.

```
[edit firewall family inet filter RE-protect term dhcp-client-accept]
user@host# set from source-address 0.0.0.0/32
user@host# set from destination-address 255.255.255.255/32
user@host# set from protocol udp
user@host# set from source-port 68
user@host# set from destination-port 67
```

4. Specify the action to take for matched packets.

```
[edit firewall family inet filter RE-protect term dhcp-client-accept]
user@host# set then count dhcp-client-accept
user@host# set then accept
```

5. Create a filter term for the server.

```
[edit firewall family inet filter RE-protect]
user@host# edit term dhcp-server-accept
```

6. Specify the match conditions for DHCP packets.

```
[edit firewall family inet filter RE-protect term dhcp-server-accept]
user@host# set from protocol udp
user@host# set from source-port [67 68]
user@host# set from destination-port [67 68]
```

7. Specify the action to take for matched packets.

```
[edit firewall family inet filter RE-protect term dhcp-server-accept]
user@host# set then count dhcp-client-accept
user@host# set then accept
```

Results

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

Copyright © 2017, Juniper Networks, Inc.

```
[edit]
user@host# show firewall
family inet {
  filter RE-protect {
    term dhcp-client-accept {
      from {
        source-address {
          0.0.0.0/32;
        }
        destination-address {
          255.255.255.255/32;
        }
        protocol udp;
        source-port 68;
        destination-port 67;
      }
      then {
        count dhcp-client-accept;
        accept;
      }
    }
    term dhcp-server-accept {
      from {
        protocol udp;
        source-port [ 67 68 ];
        destination-port [ 67 68 ];
      }
      then {
        count dhcp-server-accept;
        accept;
      }
    }
  }
}
```

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

Verification

To confirm that the Routing Engine DHCP protection filter is properly passing DHCP packets, perform these tasks:

- Verifying the DHCP Filter Operation on page 378

Verifying the DHCP Filter Operation

Purpose Verify that both counters increment as DHCP traffic passes to the Routing Engine.

Action From operational mode, enter the **show firewall family inet filter RE-protect** command.

```
user@host> show firewall family inet filter RE-protect
Filter: RE-protect
Counters:
Name                                                Bytes              Packets
```


Copyright © 2017, Juniper Networks, Inc.

```
dhcp-client-accept                                    328                 1
dhcp-server-accept                                    574                 1

user@host> show firewall family inet filter RE-protect
Filter: RE-protect
Counters:
Name                                                Bytes             Packets
dhcp-client-accept                                    660                 2
dhcp-server-accept                                   1152                 2
```

Meaning The output lists both configured counters, dhcp-client-accept and dhcp-server-accept. By issuing the command more than once, you can see that the byte and packet fields both show that traffic is being accepted and counted.

Related Documentation

- Port Number Requirements for DHCP Firewall Filters on page 379
- *Understanding Dynamic Firewall Filters*
- *Understanding Dynamic Firewall Filters*
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240

Port Number Requirements for DHCP Firewall Filters

When you configure a firewall filter to perform some action on DHCP packets at the Routing Engine, such as protecting the Routing Engine by allowing only proper DHCP packets, you must specify both port 67 (bootps) and port 68 (bootpc) for both the source and destination. The firewall filter acts at both the line cards and the Routing Engine.

This requirement applies to both DHCP local server and DHCP relay, but it applies only when DHCP is provided by the jdhcpd process. MX Series routers use jdhcpd. For DHCP relay, that means the configuration is required only at the **[edit forwarding-options dhcp-relay]** hierarchy level and not at the **[edit forwarding-options helpers bootp]** hierarchy level.

DHCP packets received on the line cards are encapsulated by jdhcpd with a new UDP header where their source and destination addresses are set to port 68 before being forwarded to the Routing Engine.

For DHCP relay and DHCP proxy, packets sent to the DHCP server from the router have both the source and destination UDP ports set to 67. The DHCP server responds using the same ports. However, when the line card receives these DHCP response packets, it changes both port numbers from 67 to 68 before passing the packets to the Routing Engine. Consequently the filter needs to accept port 67 for packets relayed from the client to the server, and port 68 for packets relayed from the server to the client.

Failure to include both port 67 and port 68 as described here results in most DHCP packets not being accepted.

Copyright © 2017, Juniper Networks, Inc.

For complete information about configuring firewall filters in general, see *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Related Documentation

- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 375
- Extended DHCP Local Server Overview on page 234
- Extended DHCP Relay Agent Overview on page 240
- *Understanding Dynamic Firewall Filters*

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 41

Monitoring and Managing DHCP

- Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings on page 381
- Clearing DHCP Bindings for Subscriber Access on page 382
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 384
- Monitoring DHCP Relay Server Responsiveness on page 384
- Verifying and Managing DHCP Local Server Configuration on page 385
- Verifying and Managing DHCP Relay Configuration on page 385

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.

  For DHCPv4:

  ```
  user@host> request dhcp server reconfigure all
  ```

  For DHCPv6:

  ```
  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 DHCP client.

  For DHCPv4:

  ```
  user@host> request dhcp server reconfigure 192.168.27.3
  ```

  For DHCPv6:

  ```
  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
  ```

Copyright © 2017, Juniper Networks, Inc.

- Specify the session ID of a DHCPv6 client.

  user@host> **request dhcpv6 server reconfigure 5**

- 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**

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- request dhcp server reconfigure on page 1526

Clearing DHCP Bindings for Subscriber Access

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
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.*
```

Related Documentation

- DHCP Auto Logout Overview on page 313
- Automatically Logging Out DHCP Clients on page 315

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**

Related Documentation

- *DHCP Options and Selective Traffic Processing Overview*
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 361

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
```


Copyright © 2017, Juniper Networks, Inc.

2. [edit forwarding-options dhcp-relay]
   user@host# set server-response-time 86,400

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- show dhcp relay statistics on page 1631
- show dhcpv6 relay statistics on page 1653

Verifying and Managing DHCP Local Server Configuration

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 routing-instance *customer routing instance*

- To display extended DHCP local server statistics:

  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:

  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 routing-instance *customer routing instance*

Related Documentation

- CLI Explorer

Verifying and Managing DHCP Relay Configuration

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*

Copyright © 2017, Juniper Networks, Inc.

- 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*

Related Documentation

- CLI Explorer

Copyright © 2017, Juniper Networks, Inc.

PART 3

Configuring IPv6 for Subscriber Management

- Migrating to IPv6 Using IPv4 and IPv6 Dual Stack on page 389
- Introduction to IPv6 Addresses on page 395
- Using NDRA to Provide IPv6 WAN Link Addressing on page 399
- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing on page 403
- Using DHCPv6 Prefix Delegation to Provide Subscriber LAN Addressing on page 405
- Using Both DHCPv6 IA_NA and DHCPv6 Prefix Delegation to Provide IPv6 WAN Link and Subscriber LAN Addressing on page 409
- Designs for IPv6 Addressing in a Subscriber Access Network on page 413
- Dual-Stack Access Models in a DHCP Network on page 419
- Dual-Stack Access Models in a PPPoE Network on page 427
- Configuring DHCPv6 Local Server on page 447
- Configuring DHCPv6 Relay Agent on page 455
- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- Best Practices for Configuring IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 473
- Configuring Dual Stack for PPPoE Access Networks in Which DHCP Is Used on page 479
- Configuring Dual Stack for PPPoE Access Networks in Which NDRA Is Used on page 483
- Configuring Address Assignment Pools for DHCPv6 on page 489
- Configuring the Dynamic Router Advertisement Protocol on page 493
- Examples: IPv4 and IPv6 Dual-Stack Designs on page 495
- Monitoring and Managing Dual Stack Subscribers on page 561
- Monitoring and Managing DHCPv6 on page 569

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 42

Migrating to IPv6 Using IPv4 and IPv6 Dual Stack

- Why Use IPv4/IPv6 Dual Stack? on page 389
- Basic Architecture of a Subscriber Access Dual-Stack Network on page 389
- Terms Used in IPv6 Subscriber Management Documentation on page 390
- IPv6 Addressing Requirements for a Subscriber Access Network on page 392

Why Use IPv4/IPv6 Dual Stack?

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.

Related Documentation

- Basic Architecture of a Subscriber Access Dual-Stack Network on page 389
- Terms Used in IPv6 Subscriber Management Documentation on page 390

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 3 on page 390 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.

Copyright © 2017, Juniper Networks, Inc.

Figure 3: IPv4 and IPv6 Dual-Stack Architecture in a Subscriber Access Network

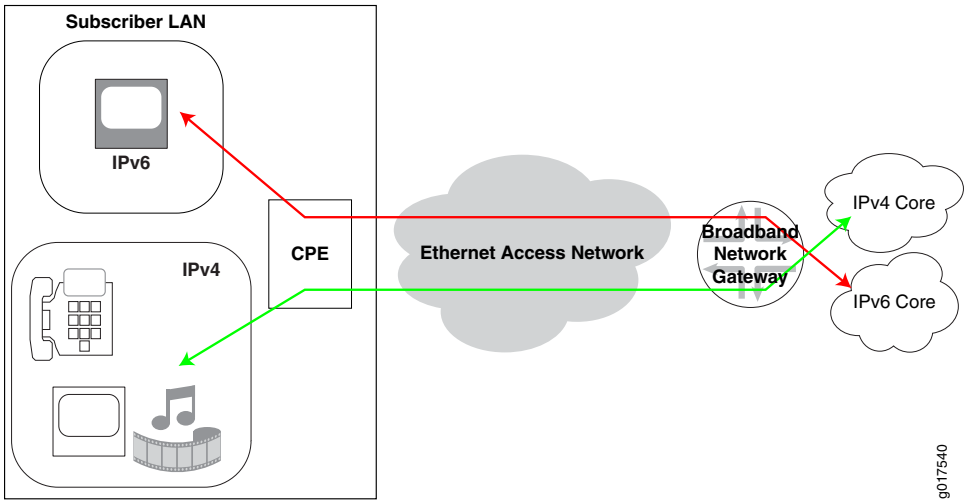

Related Documentation

- Why Use IPv4/IPv6 Dual Stack? on page 389
- Terms Used in IPv6 Subscriber Management Documentation on page 390

Terms Used in IPv6 Subscriber Management Documentation

Table 43 on page 390 defines terms used in the IPv6 subscriber management documentation.

Table 43: 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).
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.

Copyright © 2017, Juniper Networks, Inc.

Table 43: IPv6 Subscriber Management Terms *(continued)*

Term	Definition
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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 43: IPv6 Subscriber Management Terms *(continued)*

Term	Definition
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.

Related Documentation

- Why Use IPv4/IPv6 Dual Stack? on page 389
- Basic Architecture of a Subscriber Access Dual-Stack Network on page 389

IPv6 Addressing Requirements for a Subscriber Access Network

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 4 on page 392 shows where WAN link addressing and subscriber addressing are assigned in a dual-stack network.

Figure 4: IPv6 Address Requirements in a Subscriber Access Network

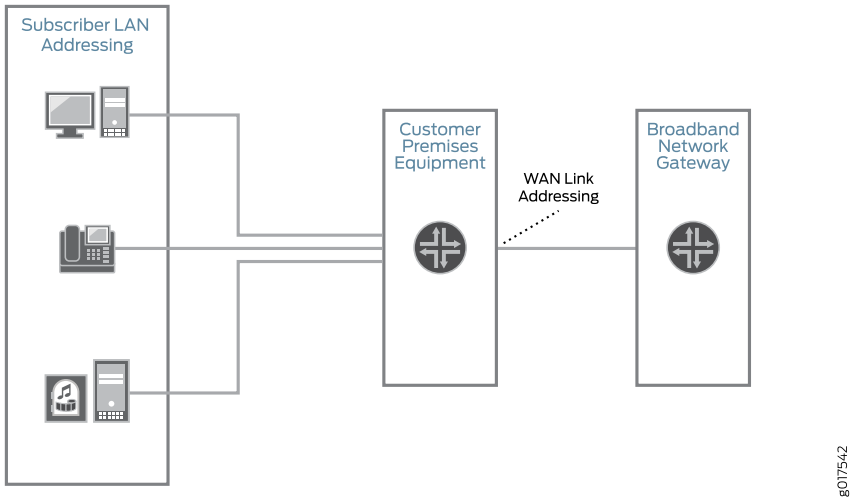

You can use the following methods for assigning IPv6 addresses:

- 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.

Copyright © 2017, Juniper Networks, Inc.

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

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403
- Using DHCPv6 Prefix Delegation Overview on page 405
- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 43

Introduction to IPv6 Addresses

- IPv6 Addressing Overview on page 395
- IPv6 Notation on page 395
- IPv6 Prefixes on page 396
- IPv6 Address Types on page 396

IPv6 Addressing Overview

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.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- IPv6 Notation on page 395
- IPv6 Prefixes on page 396
- IPv6 Address Types on page 396

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:

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- IPv6 Addressing Overview on page 395
- IPv6 Prefixes on page 396
- IPv6 Address Types on page 396
- IPv6 Addressing Requirements for a Subscriber Access Network on page 392

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.

Related Documentation

- IPv6 Addressing Overview on page 395
- IPv6 Notation on page 395
- IPv6 Address Types on page 396
- IPv6 Addressing Requirements for a Subscriber Access Network on page 392

IPv6 Address Types

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.

Copyright © 2017, Juniper Networks, Inc.

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

Copyright © 2017, Juniper Networks, Inc.

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 Overview on page 395
- IPv6 Notation on page 395
- IPv6 Prefixes on page 396
- IPv6 Addressing Requirements for a Subscriber Access Network on page 392

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 44

Using NDRA to Provide IPv6 WAN Link Addressing

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- IPv6 Neighbor Discovery Protocol Overview on page 399
- How NDRA Works in a Subscriber Access Network on page 400
- Methods for Obtaining IPv6 Prefixes for NDRA on page 401
- Duplicate Prefix Protection for NDRA on page 402

Using NDRA to Provide IPv6 WAN Link Addressing Overview

In a dual-stack network, NDRA 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.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- IPv6 Neighbor Discovery Protocol Overview on page 399
- How NDRA Works in a Subscriber Access Network on page 400
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation on page 416
- Design 3: IPv6 Addressing with NDRA on page 417
- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483

IPv6 Neighbor Discovery Protocol Overview

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.

Copyright © 2017, Juniper Networks, Inc.

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.
- 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.

Related Documentation

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- How NDRA Works in a Subscriber Access Network on page 400

How NDRA Works in a Subscriber Access Network

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).

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- IPv6 Neighbor Discovery Protocol Overview on page 399
- Methods for Obtaining IPv6 Prefixes for NDRA on page 401

Methods for Obtaining IPv6 Prefixes for NDRA

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.

Copyright © 2017, Juniper Networks, Inc.

- *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.

Related Documentation

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- How NDRA Works in a Subscriber Access Network on page 400
- Configuring an Address-Assignment Pool Used for Router Advertisements on page 486

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.

Related Documentation

- Methods for Obtaining IPv6 Prefixes for NDRA on page 401
- Configuring Duplicate IPv6 Prefix Protection for Router Advertisement on page 487

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 45

Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing

- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403
- Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA on page 403

Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview

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.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation on page 415
- Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA on page 403
- Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link on page 413
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA on page 490

Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA

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.

Copyright © 2017, Juniper Networks, Inc.

- *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.

Related Documentation

- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA on page 490

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 46

Using DHCPv6 Prefix Delegation to Provide Subscriber LAN Addressing

- Using DHCPv6 Prefix Delegation Overview on page 405
- Using a Delegated Prefix on the CPE Loopback Interface on page 407
- DHCPv6 Prefix Delegation over PPPoE on page 407
- Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation on page 408

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 5 on page 406 shows how DHCPv6 prefix delegation is used in a dual-stack network.

Copyright © 2017, Juniper Networks, Inc.

Figure 5: Delegated Addressing in a Dual-Stack Network Using DHCPv6

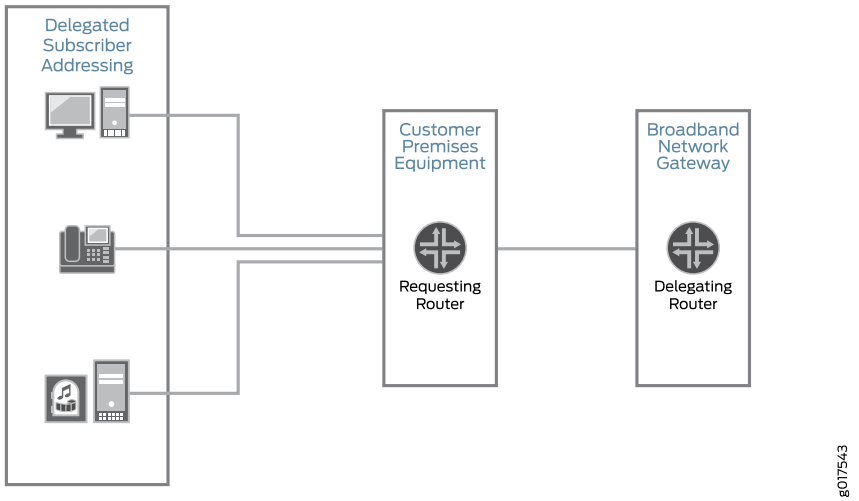

DHCPv6 prefix delegation operates as follows:

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.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation on page 415
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation on page 416
- Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix on page 418
- Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation on page 408
- Selecting the Method of Assigning Global IPv6 Addresses to Subscribers on page 414

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405
- Selecting the Type of Addressing Used on the CPE on page 413

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.

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405

Copyright © 2017, Juniper Networks, Inc.

- Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE on page 518

Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation

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.

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405
- Selecting the Method of Obtaining IPv6 Prefixes on page 414
- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation on page 489

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 47

Using Both DHCPv6 IA_NA and DHCPv6 Prefix Delegation to Provide IPv6 WAN Link and Subscriber LAN Addressing

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409
- DHCPv6 Options in a DHCPv6 Multiple Address Environment on page 410
- Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA on page 410
- Multiple DHCPv6 IA_NA and IA_PD Requests per Client Interface on page 412

Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403
- Using DHCPv6 Prefix Delegation Overview on page 405
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation on page 415
- Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE on page 495

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.**

Related Documentation

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409
- Methods for Obtaining Addresses for Both DHCPv6 Prefix Delegation and DHCPv6 IA_NA on page 410
- Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment on page 603

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

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.

Copyright © 2017, Juniper Networks, Inc.

- 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409
- DHCPv6 Options in a DHCPv6 Multiple Address Environment on page 410
- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation on page 489
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA on page 490

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.

Related Documentation

- show dhcpv6 relay binding on page 1645

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 48

Designs for IPv6 Addressing in a Subscriber Access Network

- Selecting the Type of Addressing Used on the CPE on page 413
- Selecting the Method of Provisioning a Global IPv6 Address for the WAN Link on page 413
- Selecting the Method of Assigning Global IPv6 Addresses to Subscribers on page 414
- Selecting the Method of Obtaining IPv6 Prefixes on page 414
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation on page 415
- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation on page 416
- Design 3: IPv6 Addressing with NDRA on page 417
- Design 4: IPv6 Addressing with DHCPv6 Prefix Delegation and No NDRA Prefix on page 418

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.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392

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 44 on page 414.

Copyright © 2017, Juniper Networks, Inc.

Table 44: 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.	--

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399
- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403

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.

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405
- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409

Selecting the Method of Obtaining IPv6 Prefixes

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.

Copyright © 2017, Juniper Networks, Inc.

- 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 45 on page 415 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 45: 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.
Delegated-IPv6-Prefix	DHCPv6 prefix delegation	IPv6 prefix.
Jnpr-IPv6-Delegated-Pool-Name	DHCPv6 prefix delegation	Name of an address-assignment pool configured on the BNG from which the BNG delegates a prefix.

Related Documentation

- IPv6 Addressing Requirements for a Subscriber Access Network on page 392
- Using DHCPv6 Prefix Delegation Overview on page 405

Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

This design (Figure 6 on page 416) 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.

Copyright © 2017, Juniper Networks, Inc.

Figure 6: Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

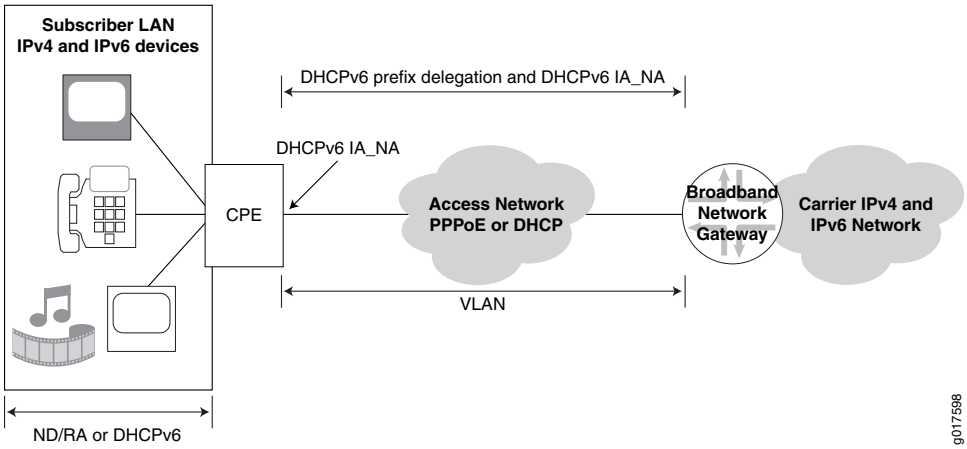

Related Documentation

- Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE on page 495
- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409

Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation

This design (Figure 7 on page 417) 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.

Copyright © 2017, Juniper Networks, Inc.

Figure 7: Subscriber Access Network with NDRA and DHCPv6 Prefix Delegation

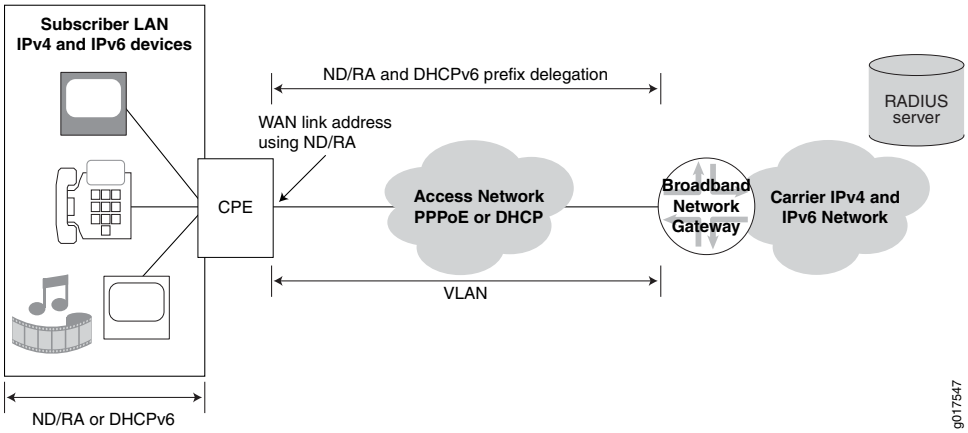

If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of design 2 and design 3.

Related Documentation

- Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE on page 518
- How NDRA Works in a Subscriber Access Network on page 400
- Using DHCPv6 Prefix Delegation Overview on page 405

Design 3: IPv6 Addressing with NDRA

In this design (Figure 8 on page 417), 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 8: Subscriber Access Network with NDRA

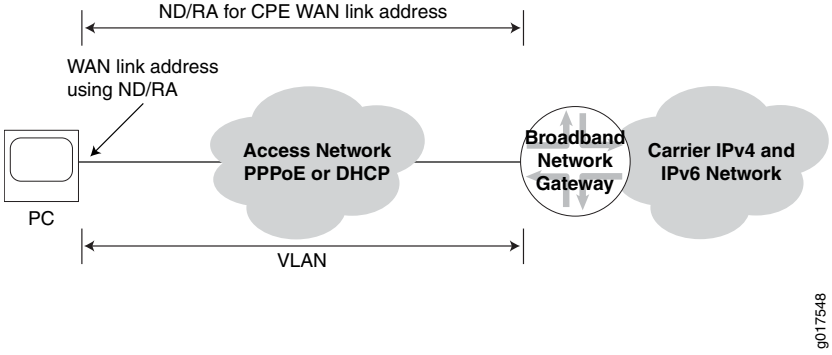

If you have a network with a combination of subscriber LANs and single PCs, you can use a combination of Design 2 and Design 3.

Related Documentation

- Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE on page 542

Copyright © 2017, Juniper Networks, Inc.

- How NDRA Works in a Subscriber Access Network on page 400

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

In this design (Figure 9 on page 418), 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 9: Subscriber Access Network with DHCPv6 Prefix Delegation

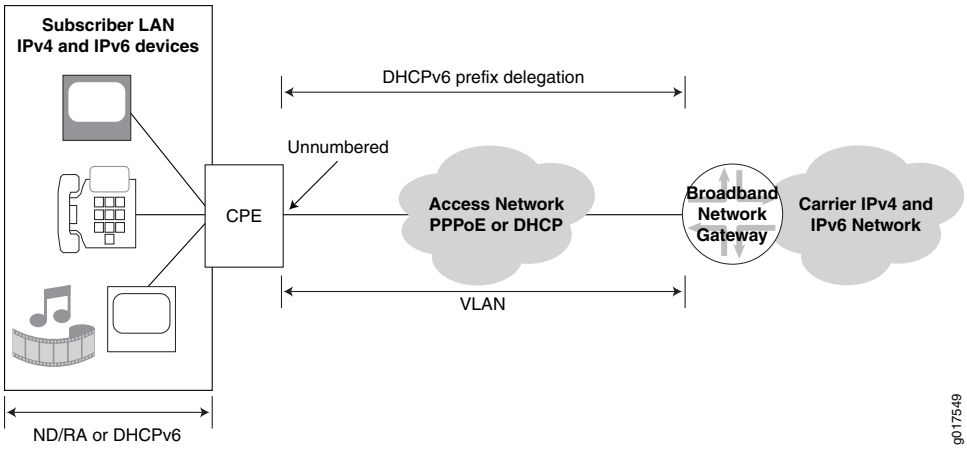

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 49

Dual-Stack Access Models in a DHCP Network

- IPv4 and IPv6 Dual Stack in a DHCP Access Network on page 419
- AAA Service Framework in a Dual Stack over a DHCP Access Network on page 420
- Dual-Stack Interface Stack in a DHCP Wholesale Network on page 421
- Single-Session DHCP Dual-Stack Overview on page 422
- Configuring Single-Session DHCP Dual-Stack Support on page 423
- Verifying and Managing DHCP Dual-Stack Configuration on page 425

IPv4 and IPv6 Dual Stack in a DHCP Access Network

Figure 10 on page 419 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 10: Dual-Stack Interface Stack over a DHCP Access Network

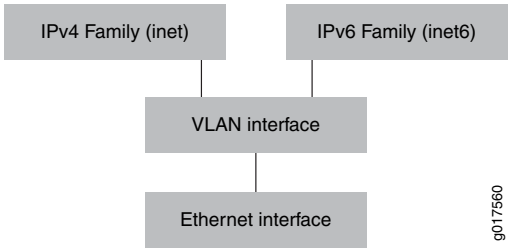

NOTE: When you are using IPv4 and IPv6 dual stack on the same DHCP interface, you must configure one dynamic profile for both the IPv4 and IPv6 subscribers. You cannot run IPv4 and IPv6 subscriber sessions over the same interface if you configure separate dynamic profiles for IPv4 and IPv6.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Basic Architecture of a Subscriber Access Dual-Stack Network on page 389
- AAA Service Framework in a Dual Stack over a DHCP Access Network on page 420
- Dual-Stack Interface Stack in a DHCP Wholesale Network on page 421

AAA Service Framework in a Dual Stack over a DHCP Access Network

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 11 on page 420, an implementation of dual stack over a DHCP access network, there are separate AAA sessions for IPv4 and IPv6 authentication and accounting.

Figure 11: AAA Service Framework in a Dual Stack over a DHCP Access Network

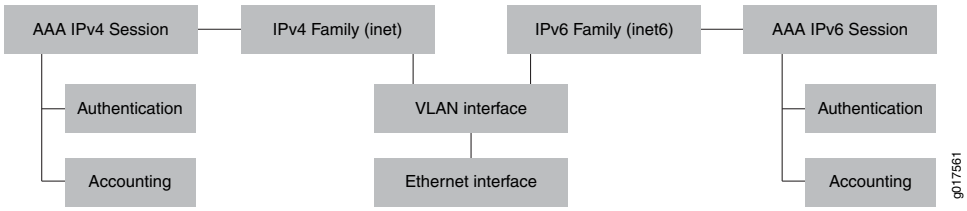

Collection of Accounting Statistics in a DHCP Access Network

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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- IPv4 and IPv6 Dual Stack in a DHCP Access Network on page 419

Dual-Stack Interface Stack in a DHCP Wholesale Network

Figure 12 on page 421 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 12: Dual-Stack Interface Stack in a DHCP Wholesale Network

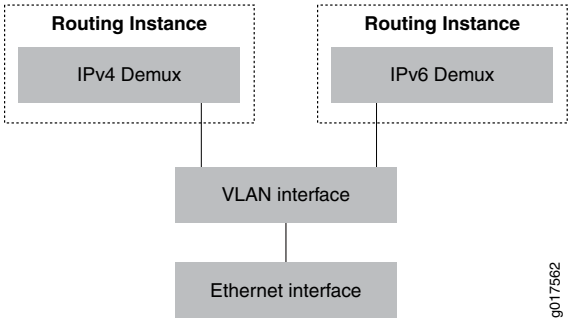

Related Documentation

- IPv4 and IPv6 Dual Stack in a DHCP Access Network on page 419
- AAA Service Framework in a Dual Stack over a DHCP Access Network on page 420

Copyright © 2017, Juniper Networks, Inc.

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 required as well, 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.

To configure single-session dual-stack subscriber settings, you use the **dual-stack-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level to create a named group that specifies the values for dual stack subscribers. Then, you use the **dual-stack** statement at the **[edit forwarding-options dhcp-relay ... overrides]** hierarchy level to specify the name of the dual stack group and assign the group to subscribers at the global, group, or interface level.

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.

- **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.
- **authentication**—Authentication-related parameters (such as password and username) the router sends to the external AAA server.

  The **dual-stack authentication** stanza is identical to the **[edit forwarding-options dhcp-relay dhcpv6]** stanza, with the added benefit that when the **username-include**

Copyright © 2017, Juniper Networks, Inc.

configuration syntax is used for the DHCPv4 leg of the dual-stack, the **relay-agent-interface-id** is equivalent to the DHCPv4 **option-82 circuit-id** statement, and the **relay-agent-remote-id** is equivalent to the DHCPv4 **option-82 remote-id** statement. You do not have to configure the two DHCPv4 options separately.

- **dynamic-profile**—Dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.
- **relay-agent-interface-id**—Includes Relay Agent Interface-ID (option 18) in DHCPv6 packets destined for the DHCPv6 server.

  For the DHCPv4 leg of the dual-stack, this statement includes the DHCPv4 **option-82 circuit-id** in packets destined for the DHCPv4 server.
- **relay-agent-remote-id**—Includes Relay Agent Remote-ID (option 37) in DHCPv6 packets destined for a DHCPv6 server.

  For the DHCPv4 leg of the dual-stack, this statement includes the DHCPv4 **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.

Related Documentation

- Configuring Single-Session DHCP Dual-Stack Support on page 423

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.

Copyright © 2017, Juniper Networks, Inc.

[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 332.

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 "Using External AAA Authentication Services with DHCP" on page 281.

   - 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 282.

   - 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 "Configuring Passwords for Usernames" on page 285.

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 458.

   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 349.

8. Specify the relay-agent-remote-id for the dual-stack group.

Copyright © 2017, Juniper Networks, Inc.

```
[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 459.

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 349.

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 324.

   - 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 324.

10. (Optional) Verify your dual-stack group configuration for DHCPv4 and DHCPv6.

    See "Verifying and Managing DHCP Dual-Stack Configuration" on page 425.

    ```
    user@host> show dhcp relay binding
    user@host> show dhcpv6 relay binding
    user@host> show subscribers
    ```

Related Documentation

- Single-Session DHCP Dual-Stack Overview on page 422

Verifying and Managing DHCP Dual-Stack Configuration

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:

Copyright © 2017, Juniper Networks, Inc.

```
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.

Related Documentation

- Single-Session DHCP Dual-Stack Overview on page 422
- Configuring Single-Session DHCP Dual-Stack Support on page 423

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 50

Dual-Stack Access Models in a PPPoE Network

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427
- Shared IPv4 and IPv6 Service Sessions on PPP Access Networks on page 430
- AAA Service Framework in a Dual Stack over a PPPoE Access Network on page 430
- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers on page 432
- Accounting Messages for PPPoE Using NDRA Prefixes on page 433
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes on page 438
- Suppressing Accounting Information That Comes from AAA on page 444
- Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address on page 444

IPv4 and IPv6 Dual Stack in a PPPoE Access Network

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 13 on page 428 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.

Copyright © 2017, Juniper Networks, Inc.

Figure 13: Dual-Stack Interface Stack over a PPPoE Access Network

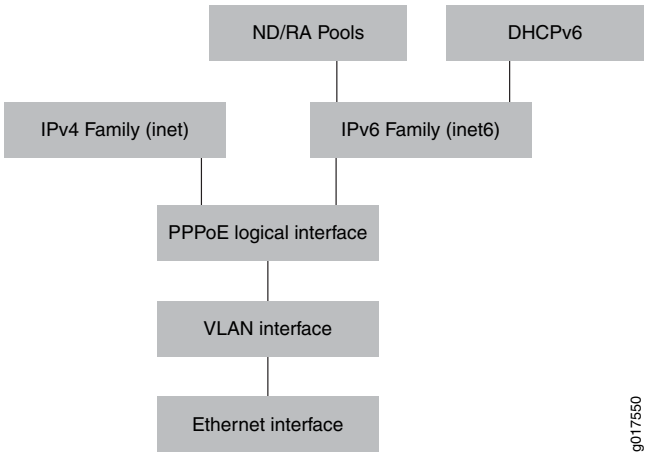

Figure 14 on page 428 shows a dual-stack interface stack over aggregated Ethernet in a PPPoE access network.

Figure 14: Dual-Stack Aggregated Ethernet Stack over a PPPoE Access Network

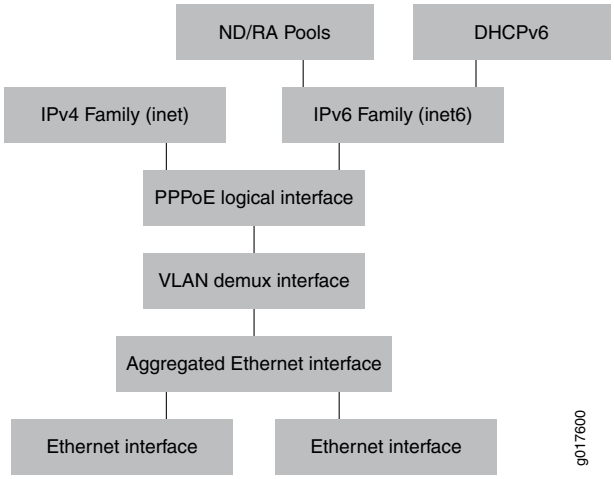

Support for Demultiplexing Interfaces

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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- AAA Service Framework in a Dual Stack over a PPPoE Access Network on page 430
- Basic Architecture of a Subscriber Access Dual-Stack Network on page 389
- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480
- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483
- Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address on page 444
- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Copyright © 2017, Juniper Networks, Inc.

Shared IPv4 and IPv6 Service Sessions on PPP Access Networks

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.

Related Documentation

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427
- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480
- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483

AAA Service Framework in a Dual Stack over a PPPoE Access Network

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.

Copyright © 2017, Juniper Networks, Inc.

As shown in Figure 15 on page 431, 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 15: AAA Service Framework in a Dual Stack over a PPPoE Access Network

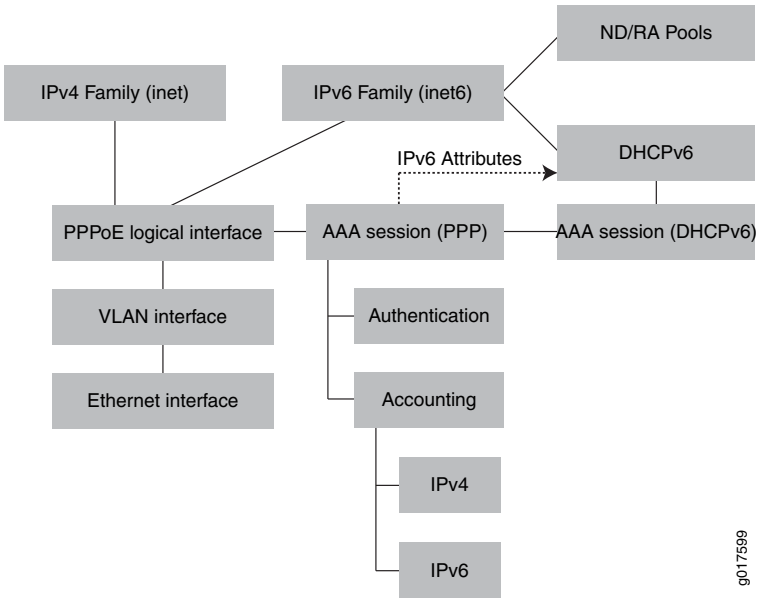

Collection of Accounting Statistics in a PPPoE Access Network

AAA provides support for both IPv4 and IPv6 statistics in one accounting session. On MX Series 3D Universal Edge routers 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427

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
2	Framed-Pool	IP Address Pool

Copyright © 2017, Juniper Networks, Inc.

Number	RADIUS Attribute	Address Type
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

Related Documentation

- Accounting Messages for PPPoE Using NDRA Prefixes on page 433
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes on page 438
- Suppressing Accounting Information That Comes from AAA on page 444

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.

Copyright © 2017, Juniper Networks, Inc.

Number	RADIUS Attributes	IPv6 Address Negotiation Type	RADIUS Accounting Messages
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.
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.

Copyright © 2017, Juniper Networks, Inc.

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.
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.

Copyright © 2017, Juniper Networks, Inc.

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.)

Copyright © 2017, Juniper Networks, Inc.

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.
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.)

Related Documentation

- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers on page 432
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes on page 438
- Suppressing Accounting Information That Comes from AAA on page 444

Copyright © 2017, Juniper Networks, Inc.

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.
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.

Copyright © 2017, Juniper Networks, Inc.

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.
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.

Copyright © 2017, Juniper Networks, Inc.

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.)
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.)

Copyright © 2017, Juniper Networks, Inc.

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.)
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.)

Copyright © 2017, Juniper Networks, Inc.

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.)

Copyright © 2017, Juniper Networks, Inc.

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. (This can occur if DHCPv6 occurs after periodic interval.)

Related Documentation

- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers on page 432
- Accounting Messages for PPPoE Using NDRA Prefixes on page 433
- Suppressing Accounting Information That Comes from AAA on page 444

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- RADIUS Accounting Messages for Dual-Stack PPPoE Subscribers on page 432
- Accounting Messages for PPPoE Using NDRA Prefixes on page 433
- Accounting Messages for PPPoE Subscribers That Use DHCPv6 IA_NA Prefixes on page 438

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,

Copyright © 2017, Juniper Networks, Inc.

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

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 51

Configuring DHCPv6 Local Server

- DHCPv6 Local Server Overview on page 448
- Enabling DHCPv6 Rapid Commit Support on page 449
- Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 450
- Configuring the DUID Type Supported by DHCPv6 Servers on page 451
- Example: Extended DHCPv6 Local Server Configuration on page 452

Copyright © 2017, Juniper Networks, Inc.

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 46 on page 448 to configure the client:

Table 46: 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.

Copyright © 2017, Juniper Networks, Inc.

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]**

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Using External AAA Authentication Services with DHCP on page 281
- Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA_NA Option on page 253
- Configuring the DUID Type Supported by DHCPv6 Servers on page 451
- Verifying and Managing DHCPv6 Local Server Configuration on page 570
- Example: Extended DHCPv6 Local Server Configuration on page 452

Enabling DHCPv6 Rapid Commit Support

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
   ```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 321
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- Extended DHCP Local Server Overview on page 234

Copyright © 2017, Juniper Networks, Inc.

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 require that all clients accept reconfiguration:

- Specify strict reconfiguration.

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

To override the global configuration for a group of clients, include the statement at the **[edit system services dhcp-local-server dhcpv6 group *group-name* reconfigure]** hierarchy level.

The **show dhcpv6 server statistics** command displays a count of solicit messages that the server has discarded.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

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;
    }
}
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 52

Configuring DHCPv6 Relay Agent

- DHCPv6 Relay Agent Overview on page 455
- DHCPv6 Relay Agent Options on page 456
- Configuring DHCPv6 Relay Agent Options on page 456
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458
- Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets on page 459

DHCPv6 Relay Agent Overview

The DHCPv6 relay agent enhances the DHCP relay agent by providing support in an IPv6 network. The DHCPv6 relay agent passes messages between the DHCPv6 client and the DHCPv6 server, similar to the way DHCP relay agent supports an IPv4 network.

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.**

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]**

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 287
- Configuring Group-Specific DHCP Relay Options on page 290
- Overriding the Default DHCP Relay Configuration Settings on page 324
- Configuring Passwords for Usernames on page 285
- Creating Unique Usernames for DHCP Clients on page 282
- Example: Extended DHCPv6 Local Server Configuration on page 452

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).

Related Documentation

- Configuring DHCPv6 Relay Agent Options on page 456
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458
- Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets on page 459

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).

Copyright © 2017, Juniper Networks, Inc.

When you enable the DHCPv6 options support, you can optionally configure DHCPv6 relay agent to include a prefix or the interface description as part of the option information. For dual-stack environments, you can also specify that the DHCPv6 relay agent use the DHCPv4 option 82 information to populate DHCPv6 option 18 or option 37.

To enable insertion of DHCPv6 options:

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

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

2. Configure DHCPv6 relay agent to insert the Relay Agent Interface-ID option, the Relay Agent Remote-ID option, or both.

   - To insert Relay Agent Interface-ID (option 18):

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# set relay-agent-interface-id
     ```

   - To insert Relay Agent Remote-ID (option 37):

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# set relay-agent-remote-id
     ```

3. (Optional) Specify additional information that you want to include in option 18 or option 37. The **relay-agent-interface-id** and **relay-agent-remote-id** statements both support inclusion of a prefix, interface description, or the DHCPv4 option 82 information. For example:

   - To prepend prefix information—This example prepends a prefix that consists of the hostname and logical system name to option 18. You use the **relay-agent-remote-id** statement to add the prefix to option 37.

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# set relay-agent-interface-id prefix host-name logical-system-name
     ```

   - To include the textual interface description—This example uses the description for the device interface instead of the interface identifier in option 18. You use the **relay-agent-remote-id** statement to add the interface description to option 37.

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# set relay-agent-interface-id use-interface-description device
     ```

   - To use the DHCPv4 option-82 value—This example uses the DHCPv4 option-82 (suboption 2) value for the DHCPv6 option 37 value. You use the **relay-agent-interface-id** statement to use DHCPv4 option 82 (suboption 1) in DHCPv6 option 18.

     This example also includes the optional **strict** keyword to specify that the router drops Solicit packets if the packets do not include an option 82 value. If you do not include the **strict** keyword, the router sends the RELAY-FORW message without adding option 37. The **strict** keyword is not supported for the relay-agent-interface-id statement.

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# set relay-agent-remote-id use-option-82 strict
     ```

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458
- Inserting DHCPv6 Remote-ID Option (Option 37) In DHCPv6 Packets on page 459
- Including a Prefix in DHCP Options on page 352
- Including a Textual Description in DHCP Options on page 354

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- DHCPv6 Relay Agent Options on page 456
- Configuring DHCPv6 Relay Agent Options on page 456
- Including a Prefix in DHCP Options on page 352
- Including a Textual Description in DHCP Options on page 354

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

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 Relay Agent Options on page 456
- Configuring DHCPv6 Relay Agent Options on page 456
- Extracting an Option 82 or Option 37 Substring to Create an Interface Set on page 358

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 53

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

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview on page 462
- On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview on page 465
- IPCP Negotiation with Optional Peer IP Address on page 468
- How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled on page 468
- Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 469
- Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470
- Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470
- Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes on page 471
- Enabling Unisphere IPv4 Release Control VSA in RADIUS Messages on page 471

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview on page 462
- On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview on page 465
- How RADIUS Attributes Are Used During Authentication When On-Demand Address Allocation is Enabled on page 468
- Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 469
- Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470
- Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470

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

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 on page 462
- IPv4 Address Release for Static PPP Subscribers on page 464

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 Unisphere-Ipv4-release-control VSA (if configured) to the RADIUS server.

Copyright © 2017, Juniper Networks, Inc.

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 Unisphere-Ipv4-release-control VSA, 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 Unisphere-Ipv4-release-control VSA, 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 Unisphere-Ipv4-release-control VSA (if configured) and the Framed-IP-Address attribute to the RADIUS server.

Copyright © 2017, Juniper Networks, Inc.

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 Unisphere-Ipv4-release-control VSA (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.

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- On-Demand IPv4 Address Negotiation and Release for Dynamic PPP Subscribers Overview on page 465

Copyright © 2017, Juniper Networks, Inc.

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

This topic describes how on-demand IPv4 address allocation and de-allocation works for static dual-stack PPP subscribers.

- IPv4 Address Negotiation for Dynamic PPP Subscribers on page 465
- IPv4 Address Release for Dynamic PPP Subscribers on page 467

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 Unisphere-Ipv4-release-control VSA (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.

Copyright © 2017, Juniper Networks, Inc.

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 Unisphere-Ipv4-release-control VSA, 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 Unisphere-Ipv4-release-control VSA, 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 Unisphere-Ipv4-release-control VSA (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.

Copyright © 2017, Juniper Networks, Inc.

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 Unisphere-Ipv4-release-control VSA (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 RAIDUS server.

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- On-Demand IPv4 Address Negotiation and Release for Static PPP Subscribers Overview on page 462

Copyright © 2017, Juniper Networks, Inc.

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.

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.

Related Documentation

- peer-ip-address-optional on page 1262
- *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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

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
   ```

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

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
   ```

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Copyright © 2017, Juniper Networks, Inc.

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

To enable the BNG to send an interim accounting messages 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
   ```

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Enabling Unisphere IPv4 Release Control VSA in RADIUS Messages

This procedure shows how to enable the Unisphere IPv4 release control VSA in RADIUS messages:

To configure a RADIUS message:

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 message RADIUS message.

   ```
   user@host# set ip-address-change-notify message
   ```

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 54

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

- Best Practice: Static PPPoE Interfaces with NDRA on page 473
- Best Practice: DHCPv6 Prefix Delegation over a PPPoE Access Network on page 474
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA on page 474
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 on page 475
- Best Practice: IPv4 Addressing for Logical Interfaces in PPPoE Dynamic Profiles on page 476
- Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network on page 477

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::1.1.1.1/64;
   }
}
```

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

Copyright © 2017, Juniper Networks, Inc.

family inet6 address configured for the interface and the prefix configured for the interface in the router advertisement configuration.

Related Documentation

- Configuring a Static PPPoE Logical Interface for NDRA on page 485

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;
```

Related Documentation

- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480

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 {
```


Copyright © 2017, Juniper Networks, Inc.

```
            address "$junos-ipv6-address ";
          }
        }
      }
    }
  }
```

Related Documentation

- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483

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;
```

Copyright © 2017, Juniper Networks, Inc.

```
        }
      }
    }
  }
}
```

Related Documentation

- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480

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;
        }
      }
    }
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
}
```

Related Documentation

- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480
- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483

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

- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE on page 479

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 55

Configuring Dual Stack for PPPoE Access Networks in Which DHCP Is Used

- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE on page 479
- Configuring a PPPoE Dynamic Profile for Use with DHCP Addressing in a Dual-Stack Network on page 480

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
   ```

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427
- Best Practice: Configuring Authentication for DHCP Subscribers on a PPPoE Access Network on page 477

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

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with DHCPv6 on page 475

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 56

Configuring Dual Stack for PPPoE Access Networks in Which NDRA Is Used

- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483
- Configuring a Static PPPoE Logical Interface for NDRA on page 485
- Configuring an Address-Assignment Pool Used for Router Advertisements on page 486
- Configuring Duplicate IPv6 Prefix Protection for Router Advertisement on page 487

Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network

Configure a dynamic profile for IPv4 and IPv6 PPPoE subscribers that access the network. The dynamic profile defines the attributes of the dynamic PPPoE logical subscriber interface.

This dynamic profile is for configurations that use NDRA to assign a global IP address to the CPE WAN link.

To configure a PPPoE dynamic profile for NDRA:

1. Create and name the dynamic profile.

   ```
   [edit]
   user@host# edit dynamic profiles PPPOE-dyn-ipv4v6-ndra
   ```

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-ndra]
     user@host# set routing-instances $junos-routing-instance interface
       $junos-interface-name
     ```

Copyright © 2017, Juniper Networks, Inc.

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-ndra]
   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 assigning an explicit IP address to the interface.

     For example:

     ```
     [edit dynamic-profiles PPPOE-dyn-ipv4v6-ndra 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-ndra 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, and assign **$junos-ipv6-address** as the predefined variable. Use this variable when you are using router advertisement with or without routing instances. This variable is replaced with the IPv6 address of the interface used for router advertisements.

   ```
   [edit dynamic profiles PPPOE-dyn-ipv4v6-ndra interfaces pp0 unit
     "$junos-interface-unit"]
   user@host# set family inet6 address $junos-ipv6-address
   ```

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-ndra pp0 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-ndra interfaces pp0 unit
     "$junos-interface-unit"]
   user@host# set pppoe-options server
   ```


Copyright © 2017, Juniper Networks, Inc.

8. (Optional) Configure the PPP authentication protocol that is used to identify and authenticate the CPE. Specify either **chap** or **pap** (or both).

```
[edit dynamic-profiles PPPOE-dyn-ipv4v6-ndra 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-ndra interfaces pp0 unit
  "$junos-interface-unit"]
user@host# set keepalives interval 30
```

10. Configure the router advertisement protocol.

    a. Access the router advertisement configuration.

```
[edit dynamic-profiles PPPOE-dyn-ipv4v6-ndra]
user@host# edit protocols router-advertisement
```

    b. Specify the interface on which the NDRA configuration is applied. Assign **$junos-interface-name** as the predefined variable. The variable is replaced with the actual name of the interface.

```
[edit dynamic-profiles PPPOE-dyn-ipv4v6-ndra protocols router-advertisement]
user@host# edit interface $junos-interface-name
```

    c. 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 PPPOE-dyn-ipv4v6-ndra protocols router-advertisement]
user@host# set prefix $junos-ipv6-ndra-prefix
```

Related Documentation

- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA on page 474
- IPv4 and IPv6 Dual Stack in a PPPoE Access Network on page 427

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
```

Copyright © 2017, Juniper Networks, Inc.

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::1.1.1.1/64
   ```

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
   ```

Related Documentation

- Best Practice: Static PPPoE Interfaces with NDRA on page 473

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- How NDRA Works in a Subscriber Access Network on page 400
- Methods for Obtaining IPv6 Prefixes for NDRA on page 401

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*

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
```

Related Documentation

- How NDRA Works in a Subscriber Access Network on page 400
- Duplicate Prefix Protection for NDRA on page 402

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 57

Configuring Address Assignment Pools for DHCPv6

- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation on page 489
- Configuring an Address-Assignment Pool for Use by DHCPv6 IA_NA on page 490
- Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation on page 490

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

This procedure shows how to configure IPv6 local addresspools 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
   ```

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405

Copyright © 2017, Juniper Networks, Inc.

- Methods for Obtaining IPv6 Prefixes for DHCPv6 Prefix Delegation on page 408

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

- Using DHCPv6 IA_NA to Provide IPv6 WAN Link Addressing Overview on page 403
- Methods for Obtaining IPv6 Global Addresses for DHCPv6 IA_NA on page 403

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: You can also use by Juniper Networks VSA 26-161 to specify the delegated address pool. The VSA-specified value takes precedence over the delegated-address statement.

To specify the pool to be used for prefix delegation:

1. Specify that you want to configure override options for DHCPv6 local server.

Copyright © 2017, Juniper Networks, Inc.

```
[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 overrides]
user@host# set delegated-pool paris-cable-12
```

Related Documentation

- Using DHCPv6 Prefix Delegation Overview on page 405
- Configuring an Address-Assignment Pool for Use by DHCPv6 Prefix Delegation on page 489

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 58

Configuring the Dynamic Router Advertisement Protocol

- Dynamic Router Advertisement Configuration Overview on page 493

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.

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring Dynamic DHCP Client Access to a Multicast Network*

Copyright © 2017, Juniper Networks, Inc.

- Configuring an Address-Assignment Pool Used for Router Advertisements on page 486

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 59

Examples: IPv4 and IPv6 Dual-Stack Designs

- Example: Configuring a Dual Stack That Uses DHCPv6 IA_NA and DHCPv6 Prefix Delegation over PPPoE on page 495
- Example: Configuring a Dual Stack That Uses ND/RA and DHCPv6 Prefix Delegation Over PPPoE on page 518
- Example: Configuring a Dual Stack That Uses ND/RA Over PPPoE on page 542

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

- Requirements on page 495
- Overview on page 495
- Configuration on page 498
- Verification on page 514

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router
- Junos OS Release 11.4 or later

Overview

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.

Copyright © 2017, Juniper Networks, Inc.

- DHCPv4 is used for subscriber LAN addressing.
- DHCPv6 subscriber sessions are layered over an underlying PPPoE subscriber session.

Topology

Figure 16: PPPoE Subscriber Access Network with DHCPv6 IA_NA and DHCPv6 Prefix Delegation

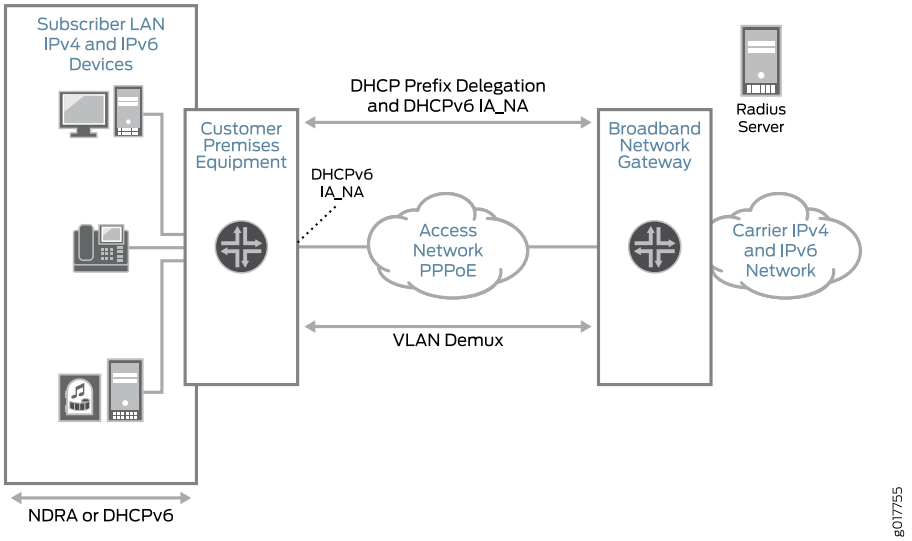

Table 47 on page 496 describes the configuration components used in this example.

Table 47: 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.
	lo0	Loopback interface for use in the access network. The loopback interface is automatically used for unnumbered interfaces.

Copyright © 2017, Juniper Networks, Inc.

Table 47: Configuration Components Used in Dual Stack with DHCPv6 IA_NA and DHCPv6 Prefix Delegation *(continued)*

Configuration Component	Component Name	Purpose
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.

Copyright © 2017, Juniper Networks, Inc.

Configuration

- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE on page 501
- Configuring a Dynamic Profile for the PPPoE Logical Interface on page 502
- Configuring a Loopback Interface on page 504
- Configuring a VLAN Demux Interface over an Ethernet Underlying Interface on page 506
- Configuring an Interface for Communication with RADIUS Server on page 508
- Specifying the BNG IP Address on page 508
- Configuring RADIUS Server Access on page 509
- Configuring RADIUS Server Access Profile on page 510
- Configuring Local Address-Assignment Pools on page 511

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;
          }
```


Copyright © 2017, Juniper Networks, Inc.

```
        }
      }
    }
  }
  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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
    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;
}
```


Copyright © 2017, Juniper Networks, Inc.

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;
                }
            }
        }
    }
}
```

Copyright © 2017, Juniper Networks, Inc.

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
    "$junos-routing-instance" {
      interface "$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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```


Copyright © 2017, Juniper Networks, Inc.

```
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

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Active Subscriber Sessions on page 514
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance on page 514
- Verifying Dynamic Subscriber Sessions on page 515
- Verifying DHCPv6 Address Pools Used for DHCPv6 Prefix Delegation on page 516
- Verifying DHCPv6 Address Bindings on page 517
- Verifying PPP Options Negotiated with the Remote Peer on page 518

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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**.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409
- Design 1: IPv6 Addressing with DHCPv6 IA_NA and DHCPv6 Prefix Delegation on page 415

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

- Requirements on page 519
- Overview on page 519

 Copyright © 2017, Juniper Networks, Inc.

- Configuration on page 521
- Verification on page 536

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router
- Junos OS Release 11.4 or later

Overview

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 17: PPPoE Subscriber Access Network with ND/RA and DHCPv6 Prefix Delegation

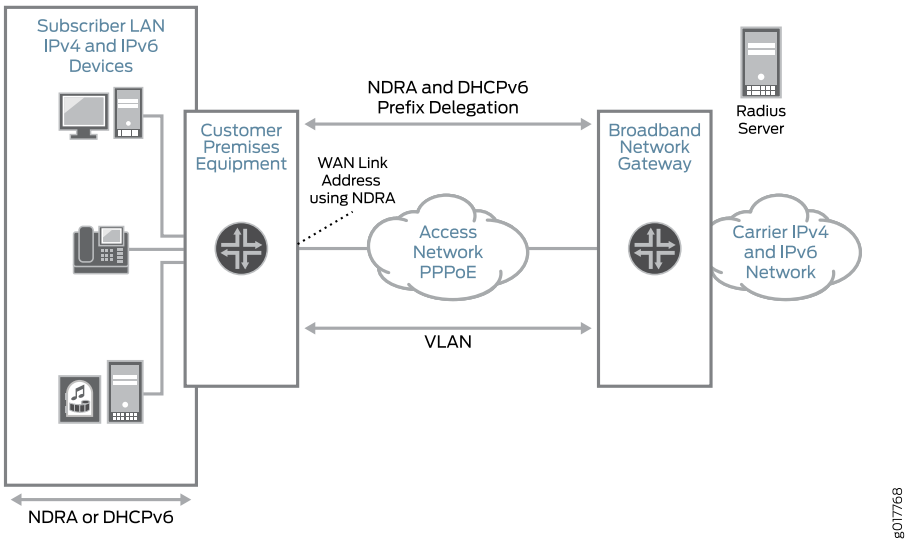

Table 48 on page 520 describes the configuration components used in this example.

Table 48: 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.
	delegated-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.

Copyright © 2017, Juniper Networks, Inc.

Configuration

- Configuring a DHCPv6 Local Server for DHCPv6 over PPPoE on page 523
- Configuring a Dynamic Profile for the PPPoE Logical Interface on page 524
- Configuring a Loopback Interface on page 527
- Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces on page 528
- Specifying the BNG IP Address on page 529
- Configuring RADIUS Server Access on page 530
- Configuring RADIUS Server Access Profile on page 531
- Specifying the RADIUS Server Access Profile to Use on page 532
- Configuring Local Address-Assignment Pools on page 533
- Specifying the Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation on page 534

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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
    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;
```


Copyright © 2017, Juniper Networks, Inc.

```
    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 delegated-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
```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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"]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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";
```

 Copyright © 2017, Juniper Networks, Inc.

```
            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]
```

Copyright © 2017, Juniper Networks, Inc.

```
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
   ```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

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:

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

```
[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 delegated-pool family inet6 prefix
  2001:db8:2040:2000:2000::/48
set address-assignment pool delegated-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 delegated-pool
user@host# edit family inet6
user@host# set prefix 2001:db8:2040:2000:2000::/48
user@host# set range r1 prefix-length 64
```

Copyright © 2017, Juniper Networks, Inc.

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 delegated-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
```


Copyright © 2017, Juniper Networks, Inc.

Step-by-Step Procedure

To specify that the dhcp-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.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Active Subscriber Sessions on page 536
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance on page 536
- Verifying Dynamic Subscriber Sessions on page 537
- Verifying DHCPv6 Address Pools Used for NDRA and DHCPv6 Prefix Delegation on page 539
- Verifying DHCPv6 Address Bindings on page 540
- Verifying Router Advertisements on page 540
- Verifying the Status of the PPPoE Logical Interface on page 541

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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: delegated-pool
IPv6 Delegated Address Pool: ndra-2010
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 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: delegated-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 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.

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Design 2: IPv6 Addressing with NDRA and DHCPv6 Prefix Delegation on page 416
- How NDRA Works in a Subscriber Access Network on page 400
- DHCPv6 Prefix Delegation over PPPoE on page 407

Copyright © 2017, Juniper Networks, Inc.

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

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 on page 542
- Overview on page 542
- Configuration on page 544
- Verification on page 557

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router
- Junos OS Release 11.4 or later

Overview

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 18: PPPoE Subscriber Access Network with NDRA

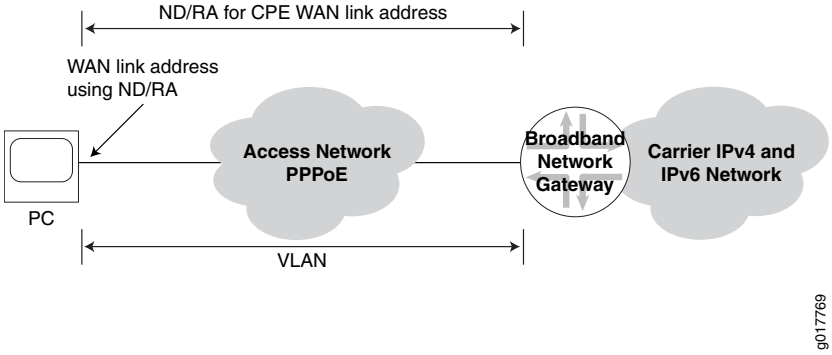

Table 49 on page 543 describes the configuration components used in this example.

Copyright © 2017, Juniper Networks, Inc.

Table 49: 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.

Copyright © 2017, Juniper Networks, Inc.

Configuration

To configure this example, perform these tasks:

- Configuring a Dynamic Profile for the PPPoE Logical Interface on page 546
- Configuring a Loopback Interface on page 549
- Configuring a Static Underlying Ethernet Interface for Dynamic PPPoE Subscriber Interfaces on page 550
- Specifying the BNG IP Address on page 551
- Configuring RADIUS Server Access on page 552
- Configuring RADIUS Server Access Profile on page 553
- Specifying the RADIUS Server Access Profile to Use on page 554
- Configuring Local Address-Assignment Pools on page 554

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 {
```


Copyright © 2017, Juniper Networks, Inc.

```
        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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
        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.

Copyright © 2017, Juniper Networks, Inc.

```
[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"]
```

Copyright © 2017, Juniper Networks, Inc.

```
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;
            }
        }
    }
```


Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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.

Copyright © 2017, Juniper Networks, Inc.

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:

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
   ```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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;
         }
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
  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.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Active Subscriber Sessions on page 557
- Verifying Both IPv4 and IPv6 Address in Correct Routing Instance on page 557
- Verifying Dynamic Subscriber Sessions on page 558
- Verifying the ND/RA Prefix Pool and Prefix Length on page 558
- Verifying the Status of the PPPoE Logical Interface on page 559
- Verifying Router Advertisements on page 560

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
```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- How NDRA Works in a Subscriber Access Network on page 400
- Design 3: IPv6 Addressing with NDRA on page 417
- Best Practice: IPv6 Addressing for Logical Interfaces in PPPoE Dynamic Profiles with NDRA on page 474

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 60

Monitoring and Managing Dual Stack Subscribers

- Monitoring Active Subscriber Sessions on page 561
- Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance on page 562
- Monitoring Dynamic Subscriber Sessions on page 562
- Monitoring Address Pools Used for Subscribers on page 563
- Monitoring Specific Subscriber Sessions on page 564
- Monitoring the Status of the PPPoE Logical Interface on page 565
- Monitoring Service Sessions for Subscribers on page 566
- Monitoring PPP Options Negotiated with the Remote Peer on page 566
- Monitoring the RADIUS Attribute Used for NDRA on page 567

Monitoring Active Subscriber Sessions

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- show subscribers summary on page 1846

Monitoring Both IPv4 and IPv6 Address in Correct Routing Instance

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.

Related Documentation
- show subscribers on page 1820

Monitoring Dynamic Subscriber Sessions

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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.

Related Documentation

- show subscribers on page 1820

Monitoring Address Pools Used for Subscribers

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: 2.2.0.5
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
```

Copyright © 2017, Juniper Networks, Inc.

```
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.

Related Documentation

- show subscribers on page 1820

Monitoring Specific Subscriber Sessions

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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.

Related Documentation

- show subscribers on page 1820

Monitoring the Status of the PPPoE Logical Interface

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- *show interfaces (PPPoE)*

Monitoring Service Sessions for Subscribers

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.

Related Documentation
- show subscribers on page 1820

Monitoring PPP Options Negotiated with the Remote Peer

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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- *show ppp interface*

Monitoring the RADIUS Attribute Used for NDRA

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

- Configuring a PPPoE Dynamic Profile for Use with NDRA in a Dual-Stack Network on page 483

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 61

Monitoring and Managing DHCPv6

- Monitoring Address Bindings on the DHCPv6 Local Server on page 569
- Verifying and Managing DHCPv6 Local Server Configuration on page 570
- Verifying and Managing DHCPv6 Relay Configuration on page 570

Monitoring Address Bindings on the DHCPv6 Local Server

Purpose Display address bindings in the client table on the DHCPv6 local server.

Action To display address bindings in the client table on the DHCPv6 local server:

```
user@host>show dhcpv6 server binding detail
user@host> show dhcpv6 server binding detail
Session Id:  6
     Client IPv6 Prefix:                    2001:db8:1111:2222::/64
     Client DUID:                          LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

     State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86308 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:
/0x00010001/0x02e159c0/0x00109400/0x0001

Session Id:  7
     Client IPv6 Prefix:                    2001:db8:1111:2222::/64
     Client DUID:                          LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

     State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86308 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
```

Copyright © 2017, Juniper Networks, Inc.

```
    Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002
```

Related Documentation

- show dhcpv6 server binding on page 1657

Verifying and Managing DHCPv6 Local Server Configuration

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

Related Documentation

- CLI Explorer

Verifying and Managing DHCPv6 Relay Configuration

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

- CLI Explorer

Copyright © 2017, Juniper Networks, Inc.

PART 4

Configuring Junos OS Enhanced Subscriber Management

- Configuring Junos OS Enhanced Subscriber Management on page 573

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 62

Configuring Junos OS Enhanced Subscriber Management

- Junos OS Enhanced Subscriber Management Overview on page 573
- Configuring Junos OS Enhanced Subscriber Management on page 575
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 578
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 579

Junos OS Enhanced Subscriber Management Overview

Junos OS enhanced subscriber management is a next-generation broadband edge software architecture for wireline subscriber management. Enhanced subscriber management enables you to take advantage of increased scaling and performance for configuring and managing dynamic interfaces and services for subscriber management.

Enhanced subscriber management delivers optimized scaling and performance for the existing dynamic subscriber management feature set. Enhanced subscriber management provides feature parity with the legacy Junos OS subscriber management feature set, with certain exceptions. For a list of these feature exceptions, see the latest *Junos OS Release Notes for MX Series 3D Universal Edge Routers* for your Junos OS software.

In order to use dynamic profiles to create and manage dynamic subscriber interfaces and services, you *must* explicitly configure and enable enhanced subscriber management. When enhanced subscriber management is enabled, it handles all subscriber-management control protocol traffic (DHCP, PPP, PPPoE, L2TP, and dynamic VLAN creation) to direct the creation of subscriber sessions and their associated dynamic interfaces.

If you are using only static network configurations and static services in a business edge environment, you do not need to enable enhanced subscriber management to configure these static topologies. When enhanced subscriber management is *not* enabled, the following client applications do not support the use of dynamic profiles, the creation of dynamic interfaces, or dynamic authentication services:

- Dynamic VLANs
- PPPoE

Copyright © 2017, Juniper Networks, Inc.

- PPP
- L2TP
- DHCP

From an operational perspective, enhanced subscriber management introduces only minimal changes to existing subscriber management configuration and verification procedures. For example, enhanced subscriber management consolidates several subscriber management components previously distributed across multiple processes into a single process. As a result, enhanced subscriber management can display consolidated information for subscriber management in a single **show** command.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

Copyright © 2017, Juniper Networks, Inc.

Configuring Junos OS Enhanced Subscriber Management

Junos OS enhanced subscriber management is a next-generation broadband edge software architecture for wireline subscriber management. With enhanced subscriber management, you can take advantage of optimized scaling and performance for configuration and management of dynamic interfaces and services for subscriber management.

You can configure enhanced subscriber management on all MX Series 3D Universal Edge Routers with Modular Port Concentrators (MPCs) installed.

NOTE: You *must* configure enhanced subscriber management in order to use dynamic profiles for creating and managing dynamic subscriber interfaces and services.

Before you begin:

- Download and install Junos OS Release 15.1R3 or later.

  See *Migration, Upgrade, and Downgrade Instructions* in the *Junos OS 15.1R3 Release Notes*. You must reboot the router after the upgrade is validated and installed.

CAUTION: Because unified in-service software upgrade (unified ISSU) is not supported when you upgrade to Junos OS Release 15.1R3, all subscriber sessions and subscriber state are lost after the upgrade.

To configure Junos OS enhanced subscriber management for the first time:

1. Configure enhanced IP network services on the router.
   a. Specify that you want to configure chassis properties for the router.

      ```
      [edit]
      user@host# edit chassis
      ```

   b. Configure enhanced IP network services.

      ```
      [edit chassis]
      user@host# set network-services enhanced-ip
      ```

2. Enable enhanced subscriber management.
   a. Specify that you want to configure global services for the router.

      ```
      [edit]
      user@host# edit system services
      ```

   b. Enable enhanced subscriber management.

      ```
      [edit system services]
      user@host# set subscriber-managementenable
      ```

Copyright © 2017, Juniper Networks, Inc.

3. Allocate shared memory for both the configuration database and enhanced subscriber management by setting a maximum size for the configuration database. Juniper recommends that you give 300 MB to the configuration database for MX240, MX480, and MX960 routers, and 100 MB for MX5, MX10, MX40, MX80 and MX104 routers. Once the size of the configuration database is set, the remaining resources are automatically allocated to enhanced subscriber management. For example:

   ```
   user@host# set system configuration-databasemax-db-size 300M
   ```

4. (Optional) Enable traffic-accounting and rate-monitoring for the given interface:

   ```
   [edit dynamic-profiles profile-name interfaces unit "$junos-interface-unit"]
   user@host# set actual-transit-statistics
   ```

5. (Optional) Enable graceful Routing Engine switchover (GRES) and nonstop active routing (NSR).

   a. Enable GRES.

      ```
      [edit chassis redundancy]
      user@host# set graceful-switchover
      ```

   b. Enable NSR (recommended if you enable GRES).

      ```
      [edit routing-options]
      user@host# set nonstop-routing
      ```

   c. Configure **commit** operations to automatically synchronize the configuration between the master Routing Engine and the standby Routing Engine.

      ```
      [edit system]
      user@host# set commit synchronize
      ```

6. Commit the configuration.

   After you commit the configuration, the software prompts you to initiate a system reboot.

7. Reboot the router software to enable enhanced subscriber management.

   a. Access operational mode.

      ```
      [edit system]
      user@host# exit
      ```

   b. Reboot the software.

      ```
      user@host> request system reboot
      ```

Example The following example shows a typical configuration to enable enhanced subscriber management.

```
[edit]
chassis {
  network-services {
    enhanced-ip;
  }
```


Copyright © 2017, Juniper Networks, Inc.

```
    redundancy {
      graceful-switchover;
    }
  }
  routing-options {
    nonstop-routing;
  }
  system {
    commit synchronize;
    configuration-database {
      max-db-size 300M;
    }
    services {
      subscriber-management {
        enable;
      }
    }
  }
```

Related Documentation

- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577
- Junos OS Enhanced Subscriber Management Overview on page 573

Verifying and Managing Junos OS Enhanced Subscriber Management

Purpose View information about class of service (CoS), routing tables, active subscribers, and the subscriber database for Junos OS enhanced subscriber management.

Action

- To display dynamic subscriber interface associations for CoS classifers, rewrite rules, and scheduler maps:

  ```
  user@host> show class-of-service interface interface-name
  ```

- To display CoS associations for a dynamic interface set:

  ```
  user@host> show class-of-service interface-set interface-set-name
  ```

- To display the mapping of CoS schedulers to forwarding classes:

  ```
  user@host> show class-of-service scheduler-map
  ```

- To display CoS traffic shaping and scheduling profiles:

  ```
  user@host> show class-of-service traffic-control-profile
  ```

- To display the active entries in the routing table:

  ```
  user@host> show route
  ```

- To display detailed information about active subscribers whose IP address matches the specified address:

Copyright © 2017, Juniper Networks, Inc.

```
user@host> show subscribers address address detail
```

- To display information about how routes are mapped to specific enhanced subscriber management interfaces:

```
user@host> show system subscriber-management route
```

- To display summary information for the subscriber management database:

```
user@host> show system subscriber-management summary
```

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Junos OS Enhanced Subscriber Management Overview on page 573

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview

Before you configure subscriber secure policy traffic mirroring, note the following:

- Subscriber secure policy mirroring runs on the radius-flow-tap service infrastructure. To configure the subscriber secure policy service, you need the same privileges that are required to configure the radius-flow-tap service.
- The subscriber secure policy feature requires some system resources while mirroring, encrypting, and sending traffic to the mediation device. For example, you might elect to use a 10-Gigabit Ethernet interface for the tunnel and mediation device if you expect the amount of traffic you plan to mirror to approach 1 Gbps of actual user data.

To configure DTCP-initiated subscriber secure policy service:

1. Configure the radius-flow-tap service support for secure subscriber policy. This support includes configuring the tunnels and optional forwarding-class information that the subscriber secure policy service uses to send mirrored traffic to the content destination device.

   See *Configuring Support for Subscriber Secure Policy Mirroring*.

2. Configure the mediation device as a user on the router. This user account allows the router to receive DTCP messages from the mediation device.

   See *Configuring the Mediation Device as a User on the Router*.

3. Configure the mediation device to provision traffic mirroring on the router.

   See *Configuring the Mediation Device to Provision Traffic Mirroring*.

4. Configure a DTCP-over-SSH connection to the mediation device.

   See *Configuring a DTCP-over-SSH Connection to the Mediation Device*.

Copyright © 2017, Juniper Networks, Inc.

5. (Optional) Enable mirroring of IPv4 multicast traffic on the router.

   See *Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic*

6. Configure SNMPv3 trap support to report mirroring information to an external device.

   See *Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring*.

You can terminate an active subscriber mirroring session at any time.

See *Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions*.

Related Documentation

- *DTCP-Initiated Subscriber Secure Policy Overview*
- *Intercept-Related Events Transmitted to the Mediation Device*

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

Before you configure subscriber secure policy traffic mirroring, note the following:

- Subscriber secure policy mirroring runs on the radius-flow-tap service infrastructure. To configure the subscriber secure policy service, you must have the same privileges that are required to configure the radius-flow-tap service.
- The subscriber secure policy feature requires some system resources while mirroring, encrypting, and sending traffic to the mediation device. For example, you might elect to use a 10-Gigabit Ethernet interface for the tunnel to the mediation device if you expect the amount of traffic you plan to mirror to approach 1 Gbps of actual user data.

To configure the subscriber secure policy service:

1. Configure radius-flow-tap service support for secure subscriber policy. This support includes optional forwarding-class information that the subscriber secure policy service uses to send mirrored traffic to the content destination device.

   See *Configuring Support for Subscriber Secure Policy Mirroring*.

2. Configure an access profile that specifies the RADIUS-related support for subscriber secure policy on the router, including a list of one or more RADIUS authentication servers. The router uses the list of specified servers for both authentication and dynamic request operations. You must also configure the RADIUS dynamic request feature, which provides the CoA message support used in-session traffic mirroring.

   See *Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring*.

3. Ensure that the following support is also configured:

   - The RADIUS record of the mirrored subscriber must include the RADIUS attributes and VSAs required for subscriber secure policy mirroring. See *RADIUS Attributes Used for Subscriber Secure Policy* for descriptions of the supported attributes used in RADIUS Accept-Accept and CoA messages.

Copyright © 2017, Juniper Networks, Inc.

- The mediation device must be configured to accept the mirrored content.

4. (Optional) Enable the mirroring of IPv4 multicast traffic on the router.

   See *Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic*.

5. (Optional) Configure SNMPv3 trap support to report mirroring-related events to the mediation device.

   See *Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring*.

To terminate an active subscriber mirroring session at any time.

See *Terminating RADIUS-Initiated Subscriber Traffic Mirroring* .

Related Documentation

- *RADIUS Attributes Used for Subscriber Secure Policy*
- *Guidelines for Configuring Subscriber Secure Policy Mirroring*
- *Intercept-Related Events Transmitted to the Mediation Device*
- *Terminating RADIUS-Initiated Subscriber Traffic Mirroring*

Copyright © 2017, Juniper Networks, Inc.

PART 5

Configuring Address-Assignment Pools for Subscriber Management

- Configuring Address-Assignment Pools for Dynamic and Static Addresses on page 583

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 63

Configuring Address-Assignment Pools for Dynamic and Static Addresses

- Address-Assignment Pools Overview on page 583
- Address-Assignment Pools Licensing Requirements on page 584
- Configuring Address-Assignment Pools on page 585
- Example: Configuring an Address-Assignment Pool on page 586
- Configuring an Address-Assignment Pool Name and Addresses on page 587
- Configuring a Named Address Range for Dynamic Address Assignment on page 588
- Configuring Address-Assignment Pool Usage Threshold Traps on page 588
- Configuring Address-Assignment Pool Linking on page 589
- Configuring Address-Assignment Pool Hold-Down on page 590
- Configuring DHCP Local Address Pool Rapid Drain on page 591
- Configuring Static Address Assignment on page 594
- Configuring Duplicate IPv4 Address Protection for AAA on page 595

Address-Assignment Pools Overview

The address-assignment pool feature supports subscriber management and DHCP management functionality by enabling you to create centralized IPv4 and IPv6 address pools independently 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.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

You can link address-assignment pools together to provide backup pools for address assignment. When no addresses are found to be available in the specified (primary) address pool, the router or switch automatically proceeds to the linked (secondary) address pool to search for an available address to allocate.

The address-assignment pool hold-down feature enables you to specify that no additional address are allocated from an existing active address-assignment pool. 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.

You can also explicitly identify that an address-assignment pool is used for ND/RA.

Related Documentation

- Configuring Address-Assignment Pools on page 585
- Address-Assignment Pools Licensing Requirements on page 584
- Example: Configuring an Address-Assignment Pool on page 586

Address-Assignment Pools Licensing Requirements

The address-assignment pool feature is part of the Junos OS Subscriber Management Feature Pack license. You must install and properly configure the license to meet the requirements for using the address-assignment pool feature.

Related Documentation

- *Junos OS Feature Licenses*

Copyright © 2017, Juniper Networks, Inc.

Configuring Address-Assignment Pools

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 587.

2. (Optional) Configure named ranges (subsets) of addresses.

   See "Configuring a Named Address Range for Dynamic Address Assignment" on page 588.

3. (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 589.

4. (Optional) Create static address bindings (IPv4 only).

   See "Configuring Static Address Assignment" on page 594.

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 590.

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 591.

7. (Optional) Configure attributes for DHCP clients.

   See "Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address" on page 254.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Address-Assignment Pools Licensing Requirements on page 584
- Example: Configuring an Address-Assignment Pool on page 586
- DHCP Attributes Overview on page 244

Example: Configuring an Address-Assignment Pool

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.

```
[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 svale6.boston.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;
         }
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
}
```

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 sval6.boston.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**.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

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
   ```

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

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
```

Related Documentation
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

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 threshold traps. An address pool has SNMP thresholds associated with it that allow the local address server to signal SNMP traps when certain conditions exist. These thresholds include high utilization

Copyright © 2017, Juniper Networks, Inc.

threshold and abated utilization threshold. If a pool's outstanding addresses exceed the high utilization threshold and the SNMP trap signaling is enabled, SNMP is notified. Likewise, when a pool's utilization drops below the abated threshold utilization threshold, SNMP is notified. When the system reaches the high utilization value, it sends warning messages. When memory usage falls to the abated utilization value, the system stops sending warning messages.

To set the usage for threshold traps:

- Specify the percentage after which the address pool usage is exceeded that an SNMP trap is generated.

```
[edit access]
user@host# edit address-assignment high-utilization 95
```

To set the abated value for the trap:

- Specify the percentage below which the address pool usage is abated that an SNMP trap is generated.

```
[edit access]
user@host# edit address-assignment abated-utilization 80
```

Related Documentation

- Address-Assignment Pools Overview on page 583

Configuring Address-Assignment Pool Linking

Address-assignment pool linking enables you to specify a secondary address pool for the router to use when the primary address-assignment pool is fully allocated. When the primary pool has no available addresses, the router automatically switches over to the linked secondary pool and begins allocating addresses from that pool. The router uses a secondary pool only when the 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. When Pool_A has no available addresses, the router switches to using Pool_B for addresses. When Pool_B is exhausted, the router switches 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.

Before you begin, configure your address pools. See “Configuring Address-Assignment Pools” on page 585.

To link an address-assignment pool to a secondary pool:

- Specify the names of the pools to be linked.

```
[edit access]
user@host# set address-assignment pool-name link secondary-pool-name
```

Copyright © 2017, Juniper Networks, Inc.

For example, the following configuration links Pool_A to Pool_B and then links Pool_B to Pool_C.

```
user@host# set address-assignment pool Pool_A link Pool_B
user@host# set address-assignment pool Pool_B link Pool_C
```

Related Documentation

- Address-Assignment Pools Overview on page 583
- Address-Assignment Pools Licensing Requirements on page 584

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
   ```

Related Documentation

- Configuring Address-Assignment Pools on page 585
- Address-Assignment Pools Overview on page 583
- Configuring DHCP Local Address Pool Rapid Drain on page 591
- Address-Assignment Pools Licensing Requirements on page 584
- Example: Configuring an Address-Assignment Pool on page 586

Copyright © 2017, Juniper Networks, Inc.

Configuring DHCP Local Address Pool Rapid Drain

Copyright © 2017, Juniper Networks, Inc.

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

Copyright © 2017, Juniper Networks, Inc.

notification with a list of subscribers still having addresses from the pool that is configured for active drain. Pool draining can then continue.

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.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pool Hold-Down on page 590

Copyright © 2017, Juniper Networks, Inc.

- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- DHCP Attributes for Address-Assignment Pools on page 249

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:00:53:90ip-address 192.168.44.12
   ```

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

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.

To configure duplicate address protection:

1. Enter the **access** configuration.

   ```
   [edit]
   user@host# edit access
   ```

2. Enable duplicate address protection.

   ```
   [edit access]
   user@host# address-protection
   ```

Release History Table

Release	Description
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.

Related Documentation

- Configuring Duplicate IPv6 Prefix Protection for Router Advertisement on page 487

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 6

Configuring DNS Addresses for Subscriber Management

- Configuring DNS Address Assignments and Session Options on page 599

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 64

Configuring DNS Address Assignments and Session Options

- DNS Name Server Address Overview on page 599
- Configuring DNS Name Server Addresses for Subscriber Management on page 600
- Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 602
- Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment on page 603
- DNS Resolver for IPv6 DNS Overview on page 604
- Configuring a DNS Server Address for IPv6 Hosts on page 604

DNS Name Server Address Overview

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring DNS Name Server Addresses for Subscriber Management on page 600
- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254

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 the DHCP topics referenced in the *Related Documentation* section. 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 599.

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**.

Copyright © 2017, Juniper Networks, Inc.

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:

1. 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
   ```

2. 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:

1. 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
   ```

2. 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
   ```

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:71bfd
```

Related Documentation

- DNS Name Server Address Overview on page 599
- DHCP Attributes for Address-Assignment Pools on page 249
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses

Starting in Junos OS Release 15.1, you can configure a router to prompt any customer premises equipment (CPE) to send the IPv4 primary or secondary DNS address options in the next configuration request if the options are not included in an initial IPCP configuration request during IPCP negotiations or if the router rejects the request. This DNS option enables the router to control IPv4 DNS address provisioning for dynamic and static, terminated PPPoE and LNS subscribers. The router includes the address options in the IPCP configuration NAK message that it sends to the CPE. The CPE then negotiates both primary and secondary IPv4 DNS addresses. Using this option avoids a situation in which the CPE does not take advantage of the DNS addresses available at the router.

To configure the router to prompt the CPE to negotiate the DNS addresses for dynamic PPPoE subscribers:

- Specify the DNS negotiation option.

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"
  ppp-options]
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for static PPPoE subscribers:

- Specify the DNS negotiation option.

```
[edit interfaces interface-name ppp-options]
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for dynamic LNS subscribers:

- Specify the DNS negotiation option.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" ppp-options]
user@host# set ipcp-suggest-dns-option
```

To configure the router to prompt the CPE to negotiate the DNS addresses for static LNS subscribers:

- Specify the DNS negotiation option.

```
[edit interfaces si-slot/pic/port unit logical-unit-number ppp-options]
user@host# set ipcp-suggest-dns-option
```


Copyright © 2017, Juniper Networks, Inc.

To configure the router to prompt the CPE to negotiate the DNS addresses for tunneled PPP subscribers with an LNS user group profile:

- Specify the DNS negotiation option.

  [edit access group-profile *profile-name* ppp-options]
  user@host# set ipcp-suggest-dns-option

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, you can configure a router to prompt any customer premises equipment (CPE) to send the IPv4 primary or secondary DNS address options in the next configuration request if the options are not included in an initial IPCP configuration request during IPCP negotiations or if the router rejects the request.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile*
- *Dynamic Profiles Overview*
- *Configuring Dynamic Authentication for PPP Subscribers*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

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

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview on page 409
- DHCPv6 Options in a DHCPv6 Multiple Address Environment on page 410

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:

- Unisphere-Ipv6-Primary-Dns
- Unisphere-Ipv6-Secondary-Dns

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.

Related Documentation

- Configuring a DNS Server Address for IPv6 Hosts on page 604
- Configuring DNS Name Server Addresses for Subscriber Management on page 600

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
   ```


Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- DNS Resolver for IPv6 DNS Overview on page 604

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 7

Configuring CLI-Based Subscriber Services

- Configuring CLI-Activated Subscriber Services on page 609
- Configuring Subscriber Services with Multiple Instances on page 617
- Monitoring and Managing Subscriber Services on page 625

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 65

Configuring CLI-Activated Subscriber Services

- CLI-Activated Subscriber Services on page 609
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 613

CLI-Activated Subscriber Services

Subscriber management enables you to use the Junos OS CLI to locally activate and deactivate dynamic subscriber services. CLI-based activation and deactivation provides local control for dynamic subscriber services that is similar to subscriber management's change of authorization (CoA) feature. CoA is considered a remote activation method because the commands, or triggers, are received from a remote server, such as a RADIUS or provisioning server. Both the CoA and CLI-based methods enable you to manage services for subscribers who are currently logged in to the network—you can activate a new service for the subscriber or deactivate a current service.

The CLI-based feature activates the specified service—you cannot use it to modify a subscriber's dynamic profile instantiation or to modify user-defined variables in a dynamic profile. You can, however, include variables that are defined for the service in the dynamic profile.

Subscriber management does not support accounting for CLI-activated subscriber services. Accounting for any service is disabled by default. Therefore when you use the CLI to activate a service, it is activated with accounting disabled, and there is no way to explicitly enable accounting for the service. CLI deactivation of a service previously activated (such as by RADIUS) has no effect on accounting for that service.

CLI-based activation and deactivation is useful in service provider networks that do not use provisioning servers or RADIUS servers to activate and deactivate subscriber services. The local control provided by the CLI-based operations enables service providers to add and remove services for existing subscribers without requiring that the subscriber log out and than log in again to complete the change. For example, a service provider might allow subscribers to log in and initially use the default service, which provides basic features. After the default service is established, the provider might then use CLI-activation to upgrade qualified subscribers to an advanced service, in addition to retaining the initial

Copyright © 2017, Juniper Networks, Inc.

service. Later, the provider can use CLI-deactivation to terminate the subscriber's advanced service session. The subscriber retains the initial service until the service is deactivated.

CLI-based activation or deactivation of a subscriber service fails if any of the following conditions exist:

- A RADIUS CoA operation or a previous CLI-based activation or deactivation is currently in progress for the subscriber. Only one dynamic request can be active for the subscriber.
- A unified in-service software upgrade (unified ISSU) operation is active.
- The specified service could not be activated or deactivated.

A CLI-based activation or deactivation of a subscriber service also fails if a PCRF has successfully activated any services for the subscriber. You must override the PCRF provisioning to be able to activate or deactivate services for such a subscriber. For more information, see "Disabling PCRF Control of a Subscriber Session" on page 885.

Related Documentation

- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 613
- Default Subscriber Service Overview on page 247

Activating and Deactivating Subscriber Services Locally with the CLI

Subscriber management enables you to use the Junos OS CLI to locally activate or deactivate dynamic subscriber services for subscribers who are currently logged in to the network. You can activate an initial service for the subscriber, provide an additional service, or deactivate the subscriber's current service.

NOTE:

A CLI-based activation or deactivation of a subscriber service fails if any of the following conditions exist:

- A RADIUS CoA operation or a previous CLI-based activation or deactivation is active for the subscriber.
- A unified in-service software upgrade (unified ISSU) operation is active.
- The specified service could not be activated or deactivated.

A CLI-based activation or deactivation of a subscriber service also fails if a PCRF has successfully activated any services for the subscriber. You must override the PCRF provisioning to be able to activate or deactivate services for such a subscriber. For more information, see "Disabling PCRF Control of a Subscriber Session" on page 885.

Copyright © 2017, Juniper Networks, Inc.

To use the CLI to activate a subscriber service:

1. (Optional) Verify the subscriber's ID, and ensure that provisioning is not enabled. To display the session IDs of all current subscribers, use the **show subscribers detail** or **show network-access aaa subscribers** command.

```
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
```

2. Activate the service for the subscriber.

```
user@host> request network-access aaa subscriber add session-id 55 service-profile
  gold-service
```

3. (Optional) Verify that the new service is activated for the subscriber. (The initial **basic-service** is also listed because it has not been deactivated.)

```
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: 00:02:15
Service name: gold-service
 Service State: SvcActive
 Session ID: 57
 Session uptime: 00:00:30
```

Copyright © 2017, Juniper Networks, Inc.

To use the CLI to deactivate a subscriber service:

1. Display the active services for the specified subscriber. The following example shows that the **basic-service** and **gold-service** are active.

```
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: 00:02:15
Service name: gold-service
 Service State: SvcActive
 Session ID: 57
 Session uptime: 00:00:30
```

2. Deactivate the service for the subscriber. The following example deletes the subscriber's **basic-service** service.

```
user@host> request network-access aaa subscriber delete session-id 55 service-profile
  basic-service
```

3. (Optional) Verify that the deleted service is no longer active for the subscriber. (The **gold-service** is still listed because it has not been deactivated.)

```
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: gold-service
 Service State: SvcActive
 Session ID: 57
 Session uptime: 00:00:30
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- CLI-Activated Subscriber Services on page 609
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 613
- Default Subscriber Service Overview on page 247

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

Subscriber management enables you to use the CLI to modify a traffic-control profile that is currently applied to existing subscribers. This feature allows you to update subscribers who are initially assigned the default traffic-control profile, which might have limited features.

TIP: You specify the default traffic-control profile with the **predefined-variable-defaults** statement and the **cos-traffic-control-profile** variable at the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy level.

There are two methods you can use to modify an traffic-control profile that is in use—global and per-subscriber. The global method modifies the traffic-control profile for all subscribers currently using the traffic-control profile. The per-subscriber method modifies the traffic-control profile for a particular subscriber—all other subscribers currently using the traffic-control profile remain unaffected.

The global and per-subscriber methods share the following characteristics:

- They modify traffic-control profiles that are currently applied to active subscribers.
- Neither method creates new traffic-control profiles; they modify existing traffic-control profiles that have been previously created using the **traffic-control-profiles** statement at the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy level.
- Modifications are transparent to the active subscribers who are using the modified profile. The modified traffic-control profile is assigned without requiring any action by the subscriber.
- Both methods are useful when updating subscribers who are initially assigned the default traffic-control profile, which might have limited features. You specify the default traffic-control profile with the **predefined-variable-defaults** statement and the **cos-traffic-control-profile** variable at the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy level.

Copyright © 2017, Juniper Networks, Inc.

NOTE: To support CLI modification of traffic-control profiles in an IPv4/IPv6 dual-stack environment, you must have the **aggregate-clients replace** statement enabled at the **[edit system services dhcp-local-server group** *group-name* **dynamic-profile** *profile-name*] hierarchy

This topic includes the following tasks:

- Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers on page 614
- Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber on page 614

Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers

To make a global modification for all current subscribers assigned a particular traffic-control profile, you change one or more parameters for the traffic-control profile and **commit** the changes.

In this example, the statement changes the shaping rate for the existing traffic-control profile named **TCP-silver**. After the change, the new shaping rate applies to all subscribers currently using **TCP-silver**.

1. Access the traffic-control profile you want to modify.

   ```
   [edit dynamic-profiles business-profile class-of-service]
   user@host# edit traffic-control-profiles TCP-silver
   ```

2. Specify the parameters that you want to modify in the traffic-control profile.

   ```
   [edit dynamic-profiles business-profile class-of-service traffic-control-profiles
     TCP-silver]
   user@host# set shaping-rate 20m
   ```

3. Commit the configuration change to update the traffic-control profile. All current subscribers using **TCP-silver** now have the new **shaping-rate**.

Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber

To make a per-subscriber modification for a specific subscriber that is currently assigned a traffic-control profile, you specify the name of the new traffic-control profile to use.

In this example, the command replaces the existing traffic-control profile with the profile named **TCP-gold**. The new traffic-control profile applies only to the subscriber identified by session ID **2551**.

- Request that the traffic-control profile named **TCP-gold** be applied to session ID 2551.

  ```
  user@host> request network-access aaa subscriber modify session-id 2551
    junos-cos-traffic-control-profile TCP-gold
  ```


Copyright © 2017, Juniper Networks, Inc.

The system then displays the status message, **Successful completion**, indicating that the modification is successful. The subscriber identified by session ID 2551 now uses the **TCP-gold** traffic-control profile.

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 66

Configuring Subscriber Services with Multiple Instances

- Subscriber Services with Multiple Instances Overview on page 617
- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622

Subscriber Services with Multiple Instances Overview

Services are activated for subscribers either at login, or by using Change of Authorization (CoA) RADIUS messages or command-line interface (CLI) requests. A subscriber can have multiple instances of the same named service, provided that each instance of the subscriber service has a different set of parameters. Support for multiple instances of a subscriber service enables you to use service parameters to customize the same service to meet different needs for a particular subscriber.

- Subscriber Service Instances and Service Parameters on page 617
- CLI Deactivation of Subscriber Services with Multiple Instances on page 618
- Subscriber Services with Multiple Instances in RADIUS Accounting Messages on page 618

Subscriber Service Instances and Service Parameters

In a subscriber access network, each subscriber has its own set of services. You can configure a specific *service instance* for a particular subscriber by specifying a *service name*, also referred to as a *service profile*, and unique service parameters for that service instance. *Service parameters* can include a combination of policy lists, filters, rate-limit profiles, class of service (CoS) profiles, and interface profiles.

For example, filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter) are considered two different instances of the same service (filter-service) because their parameters, enclosed in parentheses after the service name, are different.

Each service instance is uniquely identified by the combination of its service name and service parameters. In CoA messages, the router identifies a subscriber service by its complete activation string, which consists of the service name and, if configured, one or more service parameters in the order specified.

Copyright © 2017, Juniper Networks, Inc.

CLI Deactivation of Subscriber Services with Multiple Instances

You can use the Junos OS CLI to deactivate subscriber services with multiple instances in either of the following ways:

- Deactivate a single instance of a subscriber service by specifying the name and parameters of the service to be deactivated.

  With this feature, you can deactivate a particular instance of a subscriber service while other instances of that same service remain active. For example, assume that a subscriber identified by a particular session ID has two instances of filter-service activated: filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter). If you specify "filter-service(up-filter,down-filter)" in the **request network-access aaa subscriber delete session-id** command, the router deactivates only filter-service(up-filter,down-filter); filter-service(upstream-filter,downstream-filter) remains active.

  The ability to use both service names and service parameters to identify the particular service instance to be deactivated is analogous to the subscriber service deactivation feature in use on Juniper Networks E Series Broadband Services Routers that run JunosE Software.

- Deactivate all instances of a subscriber service by specifying only the name of the service to be deactivated, with no service parameters.

  With this feature, you can deactivate all instances of the same subscriber service with a single operational command. Using the same subscriber service example, if you specify "filter-service" in the **request network-access aaa subscriber delete session-id** command, the router deactivates both filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter).

Subscriber Services with Multiple Instances in RADIUS Accounting Messages

RADIUS Acct-Start, Interim-Acct, and Acct-Stop accounting messages include the subscriber service name and, if configured, service parameters. If RADIUS logging is enabled, the router logs all subscriber service attributes, including service names and parameters, in messages sent to and received from the RADIUS authentication server.

For example, assume that the router receives the following RADIUS Access-Accept message from the RADIUS server:

```
Jul 13 12:37:02 radius-access-accept: Activate-Service (Juniper-ERX-VSA) received:
 Tag (1) filter-service(up-filter,down-filter)
```

Table 50 on page 619 shows sample logged RADIUS Acct-Start, Interim-Acct, and Acct-Stop messages that the router sends to the RADIUS server in response to the Access-Accept message. In each of these accounting messages, the Activate-Service-Session-Name is the full activation string that includes both the service name (filter-service) and service parameters (up-filter,down-filter) to identify the service instance.

Copyright © 2017, Juniper Networks, Inc.

Table 50: Subscriber Services and Service Parameters in RADIUS Accounting Messages

RADIUS Accounting Message Type	RADIUS Accounting Message Text
Acct-Start	`Jul 13 12:37:02 radius-acct-start: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)`
Interim-Acct	`Jul 13 12:47:00 radius-acct-interim: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)`
Acct-Stop	`Jul 13 12:53:59 radius-acct-stop: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)`

Related Documentation

- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622
- Verifying and Managing Subscriber Services with Multiple Instances on page 625

Deactivating a Single Instance of a Subscriber Service with Multiple Instances

For subscriber services that have multiple instances, you can use the Junos OS CLI to deactivate a service in either of the following ways:

- Deactivate a single instance of a service by specifying the name and parameters of the service to be deactivated.
- Deactivate all instances of a service by specifying only the name of the service to be deactivated.

This topic describes how to deactivate a single instance of a subscriber service.

To use the Junos OS CLI to deactivate a single instance of a subscriber service with multiple instances:

1. Display the active services for the subscriber identified by the specified session ID.

   `user@host>` show network-access aaa subscribers session-id ***subscriber-session-id*** **detail**

   For example, the following command displays the active services for the DHCP subscriber identified by session ID 6. In this example, two instances of economy-service are active: **economy-service(up-filter,down-filter)** and **economy-service(upstrm-filter,dwnstrm-filter)**. A single instance of premium-service named **premium-service(up-filter,down-filter)** is also active.

   `user@host>` show network-access aaa subscribers session-id **6 detail**

Copyright © 2017, Juniper Networks, Inc.

```
Type: dhcp
Stripped username: fran2
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
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
  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
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

2. Deactivate the specified instance of a subscriber service by specifying its service name and parameters.

   user@host> request network-access aaa subscriber delete session-id ***subscriber-session-id*** service-profile "***profile-name*(*parameters*)"**

   For example, the following command deactivates only the instance of economy-service named economy-service(up-filter,down-filter).

   user@host> request network-access aaa subscriber delete session-id 6 service-profile "economy-service(up-filter,down-filter)"

3. (Optional) Verify that the deactivated service instance is no longer active for the subscriber.

   user@host> show network-access aaa subscribers session-id ***subscriber-session-id*** detail

Copyright © 2017, Juniper Networks, Inc.

For example, the following command displays the services still active for the DHCP subscriber identified by session ID 6. In this example, **economy-service(up-filter,down-filter)** is no longer listed because it was deactivated, but **economy-service(upstrm-filter,dwnstrm-filter)** and **premium-service(up-filter,down-filter)** are still active.

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: fran2
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: 198.51.100.13.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
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
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

Related Documentation

- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622
- Verifying and Managing Subscriber Services with Multiple Instances on page 625
- Subscriber Services with Multiple Instances Overview on page 617

Copyright © 2017, Juniper Networks, Inc.

Deactivating All Instances of a Subscriber Service with Multiple Instances

For subscriber services that have multiple instances, you can use the Junos OS CLI to deactivate a service in either of the following ways:

- Deactivate a single instance of a service by specifying the name and parameters of the service to be deactivated.
- Deactivate all instances of a service by specifying only the name of the service to be deactivated.

This topic describes how to deactivate all instances of a subscriber service.

To use the Junos OS CLI to deactivate all instances of a subscriber service with multiple instances:

1. Display the active services for the subscriber identified by the specified session ID.

   ```
   user@host> show network-access aaa subscribers session-id subscriber-session-id detail
   ```

   For example, the following command displays the active services for the DHCP subscriber identified by session ID 6. In this example, two instances of economy-service are active: **economy-service(up-filter,down-filter)** and **economy-service(upstrm-filter,dwnstrm-filter)**. A single instance of premium-service named **premium-service(up-filter,down-filter)** is also active.

   ```
   user@host> show network-access aaa subscribers session-id 6 detail
   Type: dhcp
   Stripped username: fran2
   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
   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
     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
   ```


Copyright © 2017, Juniper Networks, Inc.

```
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

2. Deactivate all instances of the specified service by specifying the service name without parameters.

   user@host> request network-access aaa subscriber delete session-id ***subscriber-session-id*** service-profile "***profile-name***"

   For example, the following command deactivates both instances of economy-service.

   user@host> request network-access aaa subscriber delete session-id 6 service-profile "economy-service"

3. (Optional) Verify that all instances of the deactivated service are no longer active for the subscriber.

   user@host> show network-access aaa subscribers session-id ***subscriber-session-id*** detail

   In the following example, only **premium-service(up-filter,down-filter)** is still active. Neither **economy-service(up-filter,down-filter)** nor **economy-service(upstrm-filter,dwnstrm-filter)** is listed because all instances of economy-service were deactivated.

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: fran2
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
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Verifying and Managing Subscriber Services with Multiple Instances on page 625
- Subscriber Services with Multiple Instances Overview on page 617

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 67

Monitoring and Managing Subscriber Services

- Verifying and Managing Subscriber Services with Multiple Instances on page 625

Verifying and Managing Subscriber Services with Multiple Instances

Purpose Display information about the active services for a subscriber identified by the specified session ID.

Action The following example displays information about the active services for the DHCP subscriber identified by session ID 6.

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: 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
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
  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
```

Copyright © 2017, Juniper Networks, Inc.

```
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: premium-service
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

Meaning If parameters are configured when a subscriber service with multiple instances is activated, the **Service name** field in the **show network-access aaa subscribers session-id** command displays both the service name and, in parentheses following the service name, the service parameters. If parameters are not configured for a particular service, the **show network-access aaa subscribers session-id** command displays only the service name. The value **SvcActive** in the **Service State** field indicates that the service is active.

In this example, two instances of economy-service are active: **economy-service(up-filter,down-filter)** and **economy-service(upstrm-filter,dwnstrm-filter)**. For **premium-service**, which is also active, the command output displays only the service name, indicating that no parameters were configured for this service.

Related Documentation

- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622

Copyright © 2017, Juniper Networks, Inc.

PART 8

Configuring ANCP and the ANCP Agent for Subscribers

- Configuring ANCP Agent Neighbors and Operations on page 629
- Configuring the ANCP Agent Traffic and CoS on page 677
- Configuring the ANCP Agent and AAA on page 687
- Monitoring and Managing ANCP for Subscriber Access on page 695

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 68

Configuring ANCP Agent Neighbors and Operations

- ANCP and the ANCP Agent Overview on page 629
- ANCP Operations in Different Network Configurations on page 638
- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649
- Associating an Access Node with Subscribers for ANCP Agent Operations on page 651
- Specifying the Interval Between ANCP Adjacency Messages on page 652
- Specifying the Maximum Number of Discovery Table Entries on page 652
- Configuring the ANCP Agent for Backward Compatibility on page 653
- Specifying How Long Processes Wait for the ANCP Agent Restart to Complete on page 654
- Configuring the ANCP Agent to Learn ANCP Partition IDs on page 654
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet on page 655

ANCP and the ANCP Agent Overview

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 on page 630
- Topology Discovery on page 630
- Subscriber Services on page 631
- ANCP Interfaces and Access Loop Circuit Identifiers on page 632
- Mapping Access Lines to Interfaces and Interface Sets on page 632
- ANCP Neighbors on page 634
- Partitions on page 635

Copyright © 2017, Juniper Networks, Inc.

- Adjacency Update Messages on page 636
- Generic Response Messages and Result Codes on page 636

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:

Copyright © 2017, Juniper Networks, Inc.

- 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*.

Copyright © 2017, Juniper Networks, Inc.

ANCP Interfaces and Access Loop Circuit Identifiers

The access loop or access line in an ANCP topology consists of the physical elements between the subscriber device (CPE) and the access node. An identifier associated with the access loop serves to identify the subscriber as well. This identifier is an alphanumeric string that actually identifies the interface on the DSLAM from which subscriber requests originate. It can be referred to by various names.

- In ANCP messages, a TLV carries the access loop circuit ID, also referred to as the access line identifier, access loop circuit identifier, or access identifier.
- DHCP discovery packets can identify the line with the Agent Circuit ID suboption in the Option 82 field.
- PPPoE discovery packets can identify the line with the Agent-Circuit-ID subattribute in the DSL Forum vendor-specific tag.

Each of these identifiers is abbreviated as ACI. When the ANCP agent receives a port management message from an access node, it uses the access loop circuit identifier contained in the message to determine which logical interface or interface set corresponds to the subscriber.

You can associate an identifier with an ANCP access line by static configuration. When you configure a logical interface by specifying the interface name at the **[edit protocols ancp interfaces]** hierarchy level, include the **access-identifier** statement to associate the access loop circuit identifier with the interface. When you configure an interface set by including the **interface-set** statement at the **[edit protocols ancp interfaces]** hierarchy level, associate the access loop circuit identifier with the interface set by including the **access-identifier** statement at the **[edit protocols ancp interfaces interface-set *interface-set-name*]** hierarchy level.

When the DHCP or PPPoE discovery packet includes an ACI, the ANCP agent can dynamically map the ACI to the subscriber interface or interface set. VLANs for the subscribers are created according to a dynamic profile; these are called agent circuit identifier-based or ACI-based dynamic VLANs.

ANCP agent support for RADIUS authentication and accounting requires that both static and dynamic ACIs must be unique across the network. No two interfaces across multiple neighbors (access nodes) can share the same identifier. The DHCP and PPPoE processes do not have information about the access node IP addresses and consequently cannot distinguish between duplicate identifiers. This situation prevents the AAA services framework from correlating a DHCP or PPPoE client session with an access line for RADIUS authentication and accounting.

Mapping Access Lines to Interfaces and Interface Sets

The ANCP agent maps the ACI for subscriber access lines to an interface or interface set to apply DSL attributes received from the access node to CoS traffic shaping for the access line. The access line mapping can be statically configured with the access-identifier statement, or dynamically derived during subscriber authentication. Static mapping always supersedes dynamic mapping.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

- 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.

Copyright © 2017, Juniper Networks, Inc.

You can also configure parameters for a specific neighbor that override global or default configurations by including any of the following statements at the **[edit protocols ancp neighbor *ip-address*]** hierarchy level:

- **adjacency-timer**—Adjust the interval between adjacency messages exchanged with this neighbor.
- **ietf-mode**—Prevent the ANCP agent from operating in a backward-compatible mode for this neighbor; for neighbors that use the current IETF implementation of ANCP.
- **maximum-discovery-table-entries**—Specify how many discovery table entries are accepted from this neighbor. Include this statement at the **[edit protocols ancp]** hierarchy level to set the number of entries globally for all neighbors.
- **pre-ietf-mode**—Enable the ANCP agent to operate in a backward-compatible mode for this neighbor; for neighbors that use the original IETF implementation of ANCP (GSMPv2) rather than the current implementation. Include this statement at the **[edit protocols ancp]** hierarchy level to operate in backward-compatible mode globally for all neighbors.

RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks, defines ANCP Version 1. ANCP was originally implemented based on General Switch Management Protocol (GSMP) version 3, sub-version 1. However, the Internet community has made so many extensions and modifications to GSMPv3 in the course of developing ANCP that ANCP is no longer interoperable with GSMPv3. Consequently, ANCP neighbors must be able to dynamically detect the version that each peer supports. A joint registry codifies the GSMP and ANCP version numbers.

When an ANCP neighbor opens adjacency negotiations, it indicates the highest version of ANCP that it supports, either 0x31 for GSMPv3 or 0x32 for ANCP Version 1. (Version 1 may also be called Version 50, referring to the decimal conversion from the hexadecimal value.) If the receiving neighbor supports that version of ANCP, it returns that value when it responds to the sending neighbors. If it does not support that version, the receiving neighbor simply drops the message.

The ANCP agent stores information about active ANCP subscribers in the Junos shared database, including DSL attributes for the access lines. this storage is persistent and is removed from the database only when you delete the interface or interface set for the access line or issue one of the following commands:

- clear ancp neighbor
- clear ancp subscriber

The persistence of the storage enables PPPoE and DHCP IP demux subscribers to be properly managed by RADIUS for authentication and accounting, with their DSL attributes, even when the ANCP connection has been temporarily terminated.

Partitions

ANCP supports the division of an access node into logical partitions. Each partition creates an adjacency with a router; each partition on an access node can form adjacencies with

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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 51 on page 637 describes the result codes that can be included in a generic response message.

Table 51: 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
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)

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648
- Triggering ANCP OAM to Test the Local Loop on page 695
- *Agent Circuit Identifier-Based Dynamic VLANs Overview*
- *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*

ANCP Operations in Different Network Configurations

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.

Copyright © 2017, Juniper Networks, Inc.

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 52 on page 639 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 52: 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 53 on page 640 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.

Copyright © 2017, Juniper Networks, Inc.

Table 53: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

- 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*.

Copyright © 2017, Juniper Networks, Inc.

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 19 on page 642). 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 19: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model)

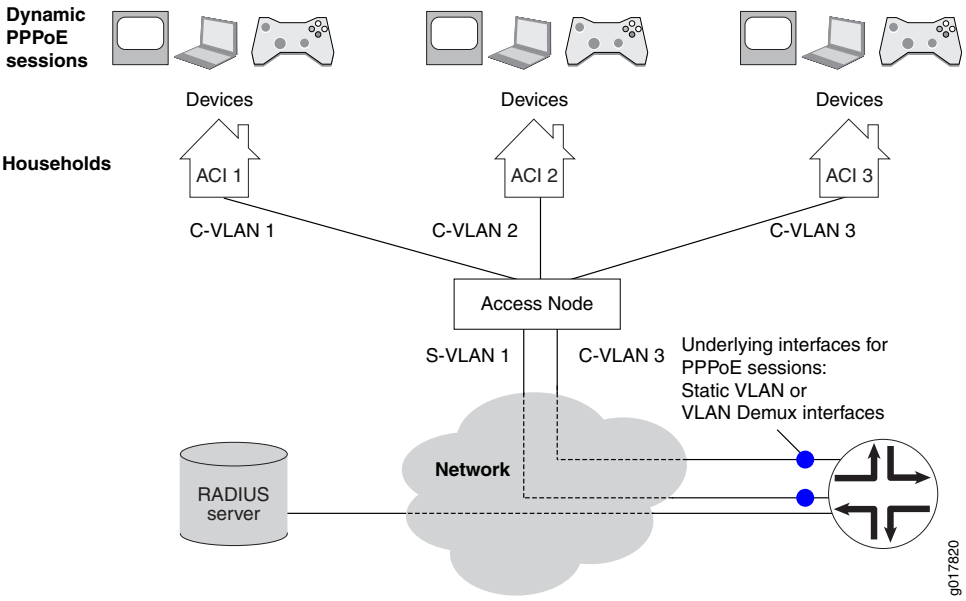

Copyright © 2017, Juniper Networks, Inc.

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 19 on page 642 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.

Copyright © 2017, Juniper Networks, Inc.

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 20 on page 644). 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 20: Sample ANCP Topology with Interface Sets (N:1 Model)

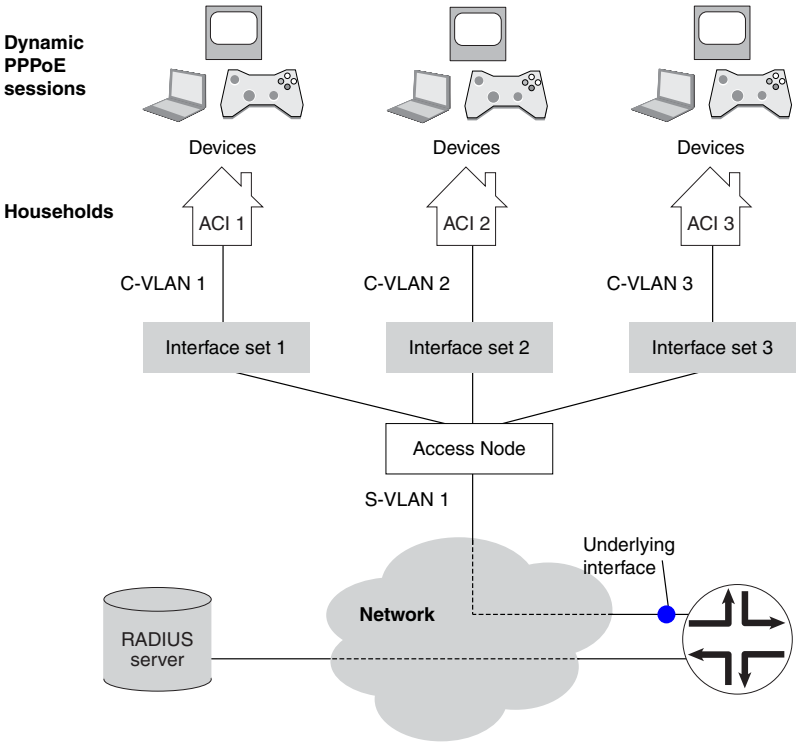

In this N:1 model with interface sets, the access node must add the DSL Forum VSA to the PPPoE PADI and PADR discovery packets that it passes to the router during the establishment of dynamic PPPoE sessions. The VSA includes the ACI for the household. This inclusion enables AAA to correlate the PPPoE 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.

Copyright © 2017, Juniper Networks, Inc.

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 20 on page 644 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

Copyright © 2017, Juniper Networks, Inc.

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 21 on page 646). 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 21: Sample ANCP Topology with Interface Sets (1:1 Model)

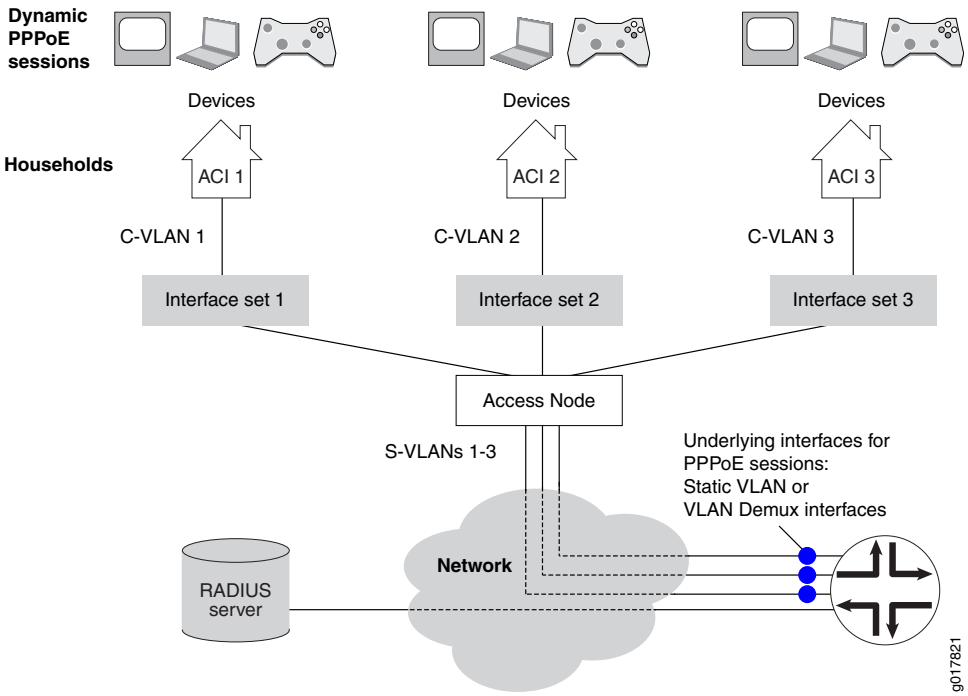

In this 1:1 model with interface sets, the ANCP agent configuration must map the underlying interface for the PPPoE sessions in an interface set to both the ACI and the interface set. This configuration enables AAA to correlate the PPPoE sessions with their respective subscriber access lines and DSL 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).

Copyright © 2017, Juniper Networks, Inc.

- 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 21 on page 646 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.

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- Configuring the ANCP Agent on page 648
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet on page 655

Copyright © 2017, Juniper Networks, Inc.

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 649.

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 651.

3. (Optional) Configure the adjacency timer.

   See "Specifying the Interval Between ANCP Adjacency Messages" on page 652.

4. (Optional) Specify the maximum number of discovery table entries that are accepted.

   See "Specifying the Maximum Number of Discovery Table Entries" on page 652

5. (Optional) Configure the ANCP agent to work with an early IETF draft.

   See "Configuring the ANCP Agent for Backward Compatibility" on page 653.

6. (Optional) Configure the graceful restart timer.

   See "Specifying How Long Processes Wait for the ANCP Agent Restart to Complete" on page 654.

7. (Optional) Configure the ANCP agent to learn partition IDs from neighbors.

   See "Configuring the ANCP Agent to Learn ANCP Partition IDs" on page 654.

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 685.

9. (Optional) Configure the ANCP agent to report unadjusted downstream traffic rates to CoS.

Copyright © 2017, Juniper Networks, Inc.

See "Configuring the ANCP Agent to Report Traffic Rates to CoS" on page 681.

10. (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 683.

11. (Optional) Configure AAA to Include or Exclude Juniper Networks DSL VSAs in RADIUS authentication and accounting messages.

    See "Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages" on page 691.

12. (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 692.

13. (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*.

14. (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*.

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

    See "Tracing ANCP Agent Operations for Subscriber Access" on page 863.

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- Triggering ANCP OAM to Test the Local Loop on page 695

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
```

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring the ANCP Agent for Backward Compatibility on page 653
- Specifying the Interval Between ANCP Adjacency Messages on page 652
- Specifying the Maximum Number of Discovery Table Entries on page 652
- *Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses*

Copyright © 2017, Juniper Networks, Inc.

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]
```

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set demux0.100 access-identifier aci_100_1_0
```

Related Documentation

- Configuring the ANCP Agent on page 648
- interfaces on page 1142

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
```

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

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
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

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
  ```

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681

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:

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols ancp]
user@host# set gsmp-syn-wait
user@host# set gsmp-syn-timeout 45
```

Related Documentation

- Configuring the ANCP Agent on page 648
- ANCP and the ANCP Agent Overview on page 629

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

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 on page 655
- Overview on page 655
- Configuration on page 660
- Verification on page 672

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Router 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 629
  - ANCP Operations in Different Network Configurations on page 638
- Configured your access node.
- Configured your RADIUS server.

Overview

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.

Copyright © 2017, Juniper Networks, Inc.

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 22 on page 656 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.

Figure 22: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet

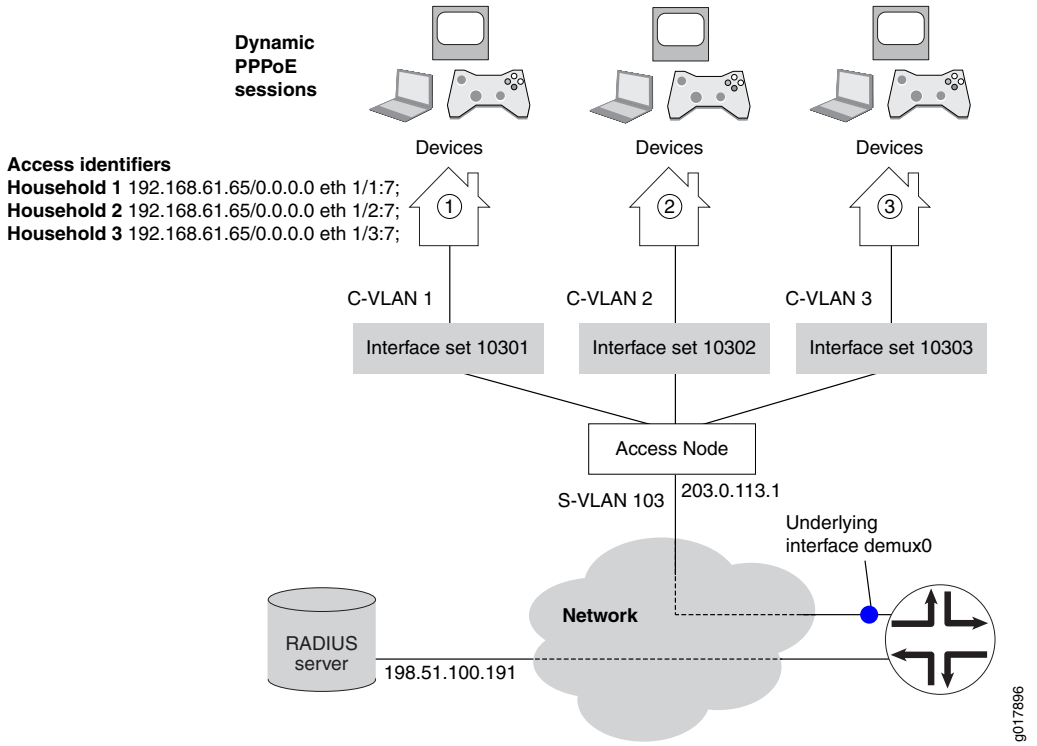

Table 54 on page 657 describes the configuration components used in this example.

Copyright © 2017, Juniper Networks, Inc.

Table 54: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 54: 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.
	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.
	pp0	PPP interface on which the PPPoE subscriber logical interfaces are created.

Copyright © 2017, Juniper Networks, Inc.

Table 54: Configuration Components used in ANCP N:1 Topology Example with Interface Sets *(continued)*

Configuration Component or Property	Component Name or Setting	Description
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.
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.

Copyright © 2017, Juniper Networks, Inc.

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

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:

- Configuring the Dynamic PPPoE Profiles on page 662
- Configuring the Static VLAN Demux Interface over Aggregated Ethernet on page 664
- Configuring Class of Service on page 669
- Configuring ANCP on page 669
- Configuring RADIUS Authentication and Accounting on page 671

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"
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
   ```


Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

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;
```

Copyright © 2017, Juniper Networks, Inc.

```
        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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

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;
```


Copyright © 2017, Juniper Networks, Inc.

```
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
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Aggregated Ethernet Interface Configuration on page 672
- Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set on page 673
- Verifying the demux0 Interface Configuration on page 673
- Verifying the pp0 Interface Configuration on page 674
- Verifying the ANCP Agent Configuration on page 675

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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:
```

Copyright © 2017, Juniper Networks, Inc.

```
    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.

Copyright © 2017, Juniper Networks, Inc.

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

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring Dynamic DHCP Client Access to a Multicast Network*
- *Subscriber Interfaces and Demultiplexing Overview*
- ANCP Agent Interactions with AAA on page 687
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 689
- Configuring the ANCP Agent on page 648
- AAA Service Framework Overview on page 3

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 69

Configuring the ANCP Agent Traffic and CoS

- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Preservation of CoS Shaping Across ANCP Agent Restarts on page 680
- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 683
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685

Traffic Rate Reporting and Adjustment by the ANCP Agent

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.

- Overview on page 677
- Traffic Rate Adjustment on page 679
- Recommended Traffic Shaping Rates on page 679
- ANCP Agent Keepalives for CoS on page 680

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. You can configure the ANCP agent to adjust the net data rate for each line type to generate the adjusted rate. The ANCP agent can adjust the rates it reports to AAA by a percentage for all DSL types. For only frame-mode DSL types (SDSL, VDSL, VDSL2, and OTHER), the agent can adjust the rates it reports to CoS by a percentage and it can adjust the number of overhead bytes per frame up or down. For cell-mode DSL types (ADSL,

Copyright © 2017, Juniper Networks, Inc.

ADSL2, and ADSL2+), the agent can adjust only the number of overhead bytes per cell up or down for traffic rates that it reports to CoS.

The ANCP agent reports traffic rates differently to AAA and CoS:

- The agent always reports both unadjusted and adjusted rates for both upstream and downstream traffic to AAA in response to a AAA request.
- The agent always 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 configured to do so, before it sends the information to CoS the agent adjusts the traffic rate for only frame-mode DSL types (SDSL, VDSL, VDSL2, and OTHER), overhead bytes for cell-mode DSL types (ADSL, ADSL2, and ADSL2+), or overhead-bytes for frame-mode DSL types (SDSL, VDSL, VDSL2, and OTHER) .

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. 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 is not always true.

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

 Copyright © 2017, Juniper Networks, Inc.

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 dynamically adjust the net data rate it reports to AAA by a fixed percentage to account for the traffic overhead. To do so, include one or more of the **qos-adjust-*dsl-line-type*** statements at the **[edit protocols ancp]** hierarchy level. Each of these statements sets an adjustment factor for a particular DSL line type such as ADSL or VDSL2. The adjustment factor is a percentage value that the ANCP agent applies to the traffic rates it receives from the DSLAM. The percentage accounts for the traffic overhead for that line type. The adjustment factor applies globally for all subscribers of the particular DSL line type associated with the statement: ADSL, ADSL2, ADSL2+, SDSL, VDSL, VDSL2, or OTHER. The ANCP agent subsequently reports the adjusted rate to AAA in addition to the unadjusted data rate.

The **qos-adjust-*dsl-line-type*** statements are enabled by default with an adjustment factor of 100 percent; by default the ANCP agent effectively makes no adjustment to the rates.

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.

For individual DSL types, you can increase the accuracy of CoS traffic shaping by configuring the agent to do one or both of the following before it reports the values to CoS:

- Adjust the net data rates by a percentage for SDSL, VDSL, VDSL2, or OTHER line types.
- Adjust the number of cell or frame overhead bytes by adding or subtracting a specified number of bytes for ADSL, ADSL2, ADSL2+, SDSL, VDSL, VDSL2, or OTHER line types.

Include the **other-overhead-adjust**, **sdsl-overhead-adjust**, **vdsl-overhead-adjust**, or **vdsl2-overhead-adjust** statements to adjust the rates. Include the **adsl-bytes**, **adsl2-bytes**, **adsl2-plus-bytes**, **other-bytes**, **sdsl-bytes**, **vdsl-bytes**, or **vdsl2-bytes** statements to adjust the bytes. These statements are all at the **[edit protocols ancp qos-adjust]** hierarchy level.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- Configuring the ANCP Agent on page 648
- *Shaping Rate Adjustments for Subscriber Local Loops Overview*
- *Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops*
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 689
- Preservation of CoS Shaping Across ANCP Agent Restarts on page 680

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- Configuring the ANCP Agent on page 648
- Specifying How Long Processes Wait for the ANCP Agent Restart to Complete on page 654

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 can 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 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). The unadjusted downstream data rate includes these headers in its calculation and reports a slightly higher value than that calculated by the DSLAM. 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 (SDSL, VDSL, VDSL2, and OTHER), 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.

Copyright © 2017, Juniper Networks, Inc.

- 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: Always configure the qos-adjust statement for normal ANCP operations. You may want to disable it for debugging purposes.

To configure the ANCP agent to report downstream traffic rates to CoS for traffic shaping:

- Enable rate reporting.

```
[edit protocols ancp]
user@host# set qos-adjust
```

To configure byte adjustment values for cell-mode DSL types (ADSL, ADSL2, or ADSL2+) or frame-mode DSL types (SDSL, VDSL, VDSL2, or OTHER):

- (Optional) Specify the number of bytes to add or subtract per cell or frame for one or more line types.

```
[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
```

You can configure also configure a rate adjustment for a frame-mode line type. Adjustments take effect immediately when they are configured.

To configure percentage adjustment values for frame-mode DSL types (SDSL, VDSL, VDSL2, or OTHER):

- (Optional) Specify the percentage for one or more line types.

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols ancp qos-adjust]
user@host# set sdsl-overhead-adjust percentage
user@host# set vdsl-overhead-adjust percentage
user@host# set vdsl2-overhead-adjust percentage;
user@host# set other-overhead-adjust percentage;
```

Related Documentation

- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648
- ANCP and the ANCP Agent Overview on page 629
- *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 on page 654
- **maximum-helper-restart-time on page 1189**

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 VSAs:

- Downstream-Calculated-Qos-Rate (IANA 4871, 26–141)—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 (IANA 4874, 26-142)—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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 691
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- ANCP and the ANCP Agent Overview on page 629

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 also reports adjusted data rates to AAA. You can configure the agent to adjust the traffic rate to account for the header overhead by including one or more of the **qos-adjust-*dsl-line-type*** statements. Each of these statements sets an adjustment factor for a particular DSL line type that applies a percentage value 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 statements 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.**

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.

```
[edit protocols ancp]
user@host# set qos-adjust-adsl adjustment-factor
user@host# set qos-adjust-adsl2 adjustment-factor
user@host# set qos-adjust-adsl2-plus adjustment-factor
user@host# set qos-adjust-other adjustment-factor
user@host# set qos-adjust-sdsl adjustment-factor
user@host# set qos-adjust-vdsl adjustment-factor
user@host# setqos-adjust-vdsl2 adjustment-factor
```

Related Documentation

- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648
- ANCP and the ANCP Agent Overview on page 629

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 70

Configuring the ANCP Agent and AAA

- ANCP Agent Interactions with AAA on page 687
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 689
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 691
- Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications on page 692

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-dsl-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.

Copyright © 2017, Juniper Networks, Inc.

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

 Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 689
- Configuring the ANCP Agent on page 648
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 691

ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes

Digital Subscriber Line (DSL) attributes are RADIUS vendor-specific attributes (VSAs) that are defined by the DSL Forum in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*. The attributes transport DSL information that is not supported by standard RADIUS attributes and which conveys details about the associated DSL subscriber line and traffic. These attributes are contained as subattributes in the single DSL Forum VSA (IANA vendor ID 3561). An ANCP access node can provide this information to the router in a PPPoE PADI message during PPPoE subscriber discovery.

The access node can also report the same information about the DSL subscriber line and traffic information by means of the ANCP DSL TLVs or attributes carried in ANCP messages to the router. The ANCP attributes are defined in RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*. These ANCP DSL attributes correspond to Juniper Networks (IANA vendor ID 4874) DSL VSAs and to DSL Forum VSAs.

The router simply passes the DSL line and traffic information that it receives from the access node to the RADIUS server, without performing any parsing or manipulation. A RADIUS authentication or accounting message can contain any combination of the DSL Forum VSAs and the Juniper Networks DSL VSAs. You can configure the RADIUS access profile to exclude one or more individual attributes, or all DSL Forum attributes, from being included in RADIUS messages.

Copyright © 2017, Juniper Networks, Inc.

The DSL Forum attribute and subattributes received by the router during PPPOE and DHCP client discovery are not updated after discovery, whereas the equivalent ANCP DSL attributes are updated whenever there is a change to the access line.

Table 55 on page 690 shows the relationship between the ANCP DSL attributes, Juniper Networks DSL VSAs, and DSL Forum VSAs.

Table 55: Mapping ANCP DSL Attributes to Juniper Networks DSL VSAs and DSL Forum VSAs

ANCP DSL Attribute Name	Juniper Networks VSA Name [Number]	DSL Forum VSA Name [Number]
Access-Aggregation-Circuit-ID-ASCII	Acc-Aggr-Cir-Id-Asc [26–112]	Access-Loop-Encapsulation [26–144]
Access-Aggregation-Circuit-ID-Binary	Acc-Aggr-Cir-Id-Bin [26–111]	Agent-Remote-Id [26–2]
Access-Loop-Circuit-ID	Acc-Loop-Cir-Id [26–110]	Agent-Circuit-Id [26–1]
Actual-Interleaving-Delay-Downstream	Act-Interlv-Delay-Dn [26–126]	Actual-Interleaving-Delay-Downstream [26–142]
Actual-Interleaving-Delay-Upstream	Act-Interlv-Delay-Up [26–124]	Actual-Interleaving-Delay-Upstream [26–140]
Actual-Net-Data-Rate-Downstream	• L2C-Down-Stream-Data [26–93]—Unadjusted rate • Act-Data-Rate-Dn [26–114]—Unadjusted rate • Downstream-Calculated-Qos-Rate [26–141]—Rate as adjusted by ANCP	Actual-Data-Rate-Downstream [26–130]
Actual-Net-Data-Rate-Upstream	• L2C-Up-Stream-Data [26–92]—Unadjusted rate • Act-Data-Rate-Up [26–113]—Unadjusted rate • Upstream-Calculated-Qos-Rate [26–142]—Rate as adjusted by ANCP	Actual-Data-Rate-Upstream [26–129]
Attainable-Net-Data-Rate-Downstream	Att-Data-Rate-Dn [26–118]	Attainable-Data-Rate-Downstream [26–134]
Attainable-Net-Data-Rate-Upstream	Att-Data-Rate-Up [26–117]	Attainable-Data-Rate-Upstream [26–133]
DSL-Line-State	DSL-Line-State [26–127]	–
DSL-Type	DSL-Type [26–128]	–
Maximum-Net-Data-Rate-Downstream	Max-Data-Rate-Dn [26–120]	Maximum-Data-Rate-Downstream [26–136]
Maximum-Net-Data-Rate-Upstream	Max-Data-Rate-Up [26–119]	Maximum-Data-Rate-Upstream [26–135]

Copyright © 2017, Juniper Networks, Inc.

Table 55: Mapping ANCP DSL Attributes to Juniper Networks DSL VSAs and DSL Forum VSAs *(continued)*

ANCP DSL Attribute Name	Juniper Networks VSA Name [Number]	DSL Forum VSA Name [Number]
Maximum-Interleaving-Delay-Downstream	Max-Interlv-Delay-Dn [26–125]	Maximum-Interleaving-Delay-Downstream [26–141]
Maximum-Interleaving-Delay-Upstream	Max-Interlv-Delay-Up [26–123]	Maximum-Interleaving-Delay-Upstream [26–139]
Minimum-Net-Low-Power-Data-Rate-Downstream	Min-LP-Data-Rate-Dn [26–122]	Minimum-Data-Rate-Downstream-Low-Power [26–138]
Minimum-Net-Low-Power-Data-Rate-Upstream	Min-LP-Data-Rate-Up [26–121]	Minimum-Data-Rate-Upstream-Low-Power [26–137]
Minimum-Net-Data-Rate-Downstream	Min-Data-Rate-Dn [26–116]	Minimum-Data-Rate-Downstream [26–132]
Minimum-Net-Data-Rate-Upstream	Min-Data-Rate-Up [26–115]	Minimum-Data-Rate-Upstream [26–131]

Related Documentation

- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 691

Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages

You can include the **juniper-dsl-attributes** statement to configure AAA to add the set of Juniper Networks DSL VSAs to the RADIUS authentication and accounting request messages for subscribers. By default, these VSAs are not added to any RADIUS message. See "ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes" on page 689 for a table of the Juniper Networks DSL VSAs.

After you have configured the inclusion of the Juniper Networks 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 DSL 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 56 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 DSL VSAs to RADIUS messages:

- Configure the inclusion trigger.

Copyright © 2017, Juniper Networks, Inc.

```
[edit access profile profile-name radius options]
user@host# set juniper-dsl-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
  ```

Related Documentation

- Configuring the ANCP Agent on page 648
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 683
- Juniper Networks VSAs Supported by the AAA Service Framework on page 21
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 689
- ANCP and the ANCP Agent Overview on page 629

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.

Copyright © 2017, Juniper Networks, Inc.

```
[edit access profile profile-name accounting]
user@host# set ancp-speed-change-immediate-update
```

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring Per-Subscriber Session Accounting on page 105

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 71

Monitoring and Managing ANCP for Subscriber Access

- Triggering ANCP OAM to Test the Local Loop on page 695
- Verifying and Monitoring ANCP Neighbors on page 696
- Clearing ANCP Neighbors on page 697
- Verifying and Monitoring ANCP Subscribers on page 697
- Clearing ANCP Subscribers on page 698
- Verifying and Monitoring CoS for ANCP Subscribers on page 698
- Clearing and Verifying ANCP Statistics on page 699

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**

Copyright © 2017, Juniper Networks, Inc.

- 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**

Related Documentation

- ANCP and the ANCP Agent Overview on page 629
- Configuring the ANCP Agent on page 648

Verifying and Monitoring ANCP Neighbors

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

Related Documentation

- CLI Explorer

Copyright © 2017, Juniper Networks, Inc.

Clearing ANCP Neighbors

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**

Related Documentation

- CLI Explorer

Verifying and Monitoring ANCP Subscribers

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

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- CLI Explorer

Clearing ANCP Subscribers

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

Related Documentation
- CLI Explorer

Verifying and Monitoring CoS for ANCP Subscribers

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**

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- CLI Explorer

Clearing and Verifying ANCP Statistics

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 and the ANCP Agent Overview on page 629

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 9

Configuring the Diameter Base Protocol and its Applications

- Configuring the Diameter Base Protocol on page 703
- Configuring Gx-Plus for Provisioning Subscribers on page 729
- Configuring 3GPP Policy and Charging Control on page 741
- Configuring NASREQ for Authentication and Authorization on page 779
- Configuring JSRC in Subscriber Access Networks on page 783
- Excluding Diameter AVPs from JSRC Messages on page 791
- Configuring Service Accounting with JSRC on page 793
- Configuring Subscribers on Static Interfaces on page 797
- Configuring the Static Subscribers Global Profile on page 805
- Configuring the Static Subscribers Group Profile on page 809
- Monitoring and Managing Diameter Information for Subscriber Access on page 815
- Monitoring and Managing Subscriber Information on Static Interfaces on page 819

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 72

Configuring the Diameter Base Protocol

- Diameter Base Protocol Overview on page 703
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Configuring Diameter on page 723
- Configuring the Origin Attributes of the Diameter Instance on page 724
- Configuring Diameter Peers on page 724
- Configuring the Diameter Transport on page 725
- Configuring Diameter Network Elements on page 726

Diameter Base Protocol Overview

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 each runs in a different Diameter instance. The individual application provides the extended AAA functionality. Applications that use Diameter include Gx-Plus, JSRC, NASREQ, and PTSP. 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, a function partition, and a metric. When an application sends a message to a routed destination, all routes within the Diameter 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

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 *master* routing instance. 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.

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Related Documentation

- Configuring Diameter on page 723
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Gx-Plus for Provisioning Subscribers Overview on page 729

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 56 on page 707 describes Diameter messages the applications use.

Copyright © 2017, Juniper Networks, Inc.

Table 56: 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).

Copyright © 2017, Juniper Networks, Inc.

Table 56: 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. 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 statistics. The PCRF may return a CCA-U message that includes new monitoring thresholds, service activations, service deactivations. A CCR-U is also sent to report the status of service activation or deactivation. 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.

Copyright © 2017, Juniper Networks, Inc.

Table 56: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 56: 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.
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.
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.

Release History Table

Release	Description
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Understanding JSRC-SAE Interactions on page 785
- Gx-Plus for Provisioning Subscribers Overview on page 729
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 731

Copyright © 2017, Juniper Networks, Inc.

Diameter AVPs and Diameter Applications

Diameter conveys information by including various attribute-value pairs (AVPs) in Diameter messages. Table 57 on page 711 lists the standard Diameter AVPs used in interactions with the supported Diameter applications. Diameter reserves AVP code numbers 0 through 255 for RADIUS AVPs that are implemented in Diameter. Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Table 57: 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
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 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

Copyright © 2017, Juniper Networks, Inc.

Table 57: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 57: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 57: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
264	Origin-Host	NASREQ	Specifies the host that originates a Diameter message.	DiamIdent

Copyright © 2017, Juniper Networks, Inc.

Table 57: 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: • 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.	Unsigned32

Copyright © 2017, Juniper Networks, Inc.

Table 57: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 57: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 57: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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 an enterprise number of 2636. Table 58 on page 718 lists the Juniper Networks AVPs that the supported Diameter applications use.

Table 58: Juniper Networks Diameter AVPs

Attribute Number	Diameter AVP	Application	Description	Type
2004	Juniper-Service-Bundle	JSRC	Specifies the name of the service bundle.	OctetString
2010	Juniper-DHCP-Options	JSRC	Specifies the client's DHCP options.	OctetString
2011	Juniper-DHCP-GI-Address	JSRC	Specifies the DHCP relay agent's IP address.	OctetString
2020	Juniper-Policy-Install	JSRC, PTSP	Specifies policies to be activated for the subscriber. Includes Juniper-Policy-Name and Juniper-Policy-Definition	Grouped
2021	Juniper-Policy-Name	JSRC, PTSP	Defines the name of a policy decision.	OctetString
2022	Juniper-Policy-Definition	JSRC, PTSP	Defines a policy decision. Includes Juniper-Policy-Name, Juniper-Template-Name, and Juniper-Substitution.	Grouped

Copyright © 2017, Juniper Networks, Inc.

Table 58: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
2023	Juniper-Template-Name	JSRC, PTSP	Specifies the profile name defined by the router. PTSP supports only the **__svc_rule__** policy template.	UTF8String
2024	Juniper-Substitution	JSRC, PTSP	Defines the substitution attributes. Includes Juniper-Substitution-Name and Juniper-Substitution-Value.	OctetString
2025	Juniper-Substitution-Name	JSRC, PTSP	Defines the name of the variable to be replaced.	OctetString
2026	Juniper-Substitution-Value	JSRC, PTSP	Defines the value of the variable to be replaced.	OctetString
2027	Juniper-Policy-Remove	JSRC, PTSP	Specifies policies to be deactivated for the subscriber. Includes Juniper-Policy-Name.	Grouped
2035	Juniper-Policy-Failed	JSRC, PTSP	Specifies the name of the policy activation or deactivation that failed.	OctetString
2038	Juniper-Policy-Success	JSRC, PTSP	Specifies the name of the policy activation or deactivation that succeeded.	OctetString
2046	Juniper-Logical-System	JSRC, PTSP	Specifies the logical system.	UTF8String
2047	Juniper-Routing-Instance	JSRC, PTSP	Specifies the routing instance.	UTF8String
2048	Juniper-Jsrc-Partition	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	JSRC, PTSP	Describes the type of request: • 1—ADDRESS_AUTHORIZATION • 2—PROVISIONING_REQUEST • 3—SYNCHRONIZATION	Enumerated
2051	Juniper-Synchronization-Type	JSRC, PTSP	Describes the type of synchronization: • 1—FULL-SYNC • 2—FAST-SYNC • 3—NO-STATE-TO-SYNC	Enumerated
2052	Juniper-Synchronization	JSRC, PTSP	Describes the state of synchronization: • 1—NO-SYNC; this is the default state • 2—SYNC-IN-PROGRESS • 3—SYNC-COMPLETE	Enumerated

Copyright © 2017, Juniper Networks, Inc.

Table 58: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
2053	Juniper-Acct-Record	JSRC, PTSP	Specifies the statistics data for each policy installed for this subscriber. Includes Juniper-Policy-Name.	Grouped
2054	Juniper-Acct-Collect	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	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	Gx-Plus, JSRC	Specifies the name of the virtual router associated with the session.	UTF8String
2101	Juniper-Provisioning-Source	Gx-Plus	Specifies the provisioning source for the session in CCR-N and JSDA messages: • 1—Local • 2—Remote	Enumerated
2102	Juniper-Provisioning-Descriptor	Gx-Plus	Defines the group used in JSDA messages that includes the session ID, and optionally Juniper-Provisioning-Source and subscriber data.	Grouped

Copyright © 2017, Juniper Networks, Inc.

Table 58: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
2103	Juniper-Event-Type	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
2104	Juniper-Discovery-Descriptor	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	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	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	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String
2108	Juniper-String-B	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String

Copyright © 2017, Juniper Networks, Inc.

Table 58: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
2109	Juniper-String-C	Gx-Plus	Specifies a generic string that is interpreted according to the context.	UTF8String
2110	Juniper-Unsigned32-A	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32
2111	Juniper-Unsigned32-B	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32
2112	Juniper-Unsigned32-C	Gx-Plus	Specifies a generic, unsigned 32-bit integer that is interpreted according to the context.	Unsigned32

Tekelec AVPs are used only for Gx-Plus. These AVPs have an enterprise number of 21274. Table 59 on page 722 lists the Tekelec AVPs. These four variables are used to provide substitution values for user-defined CoS service variables.

Table 59: 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
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

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
13.1R1	Starting in Junos OS Release 13.1R1, the packet-triggered subscribers and policy control (PTSP) feature is no longer supported.

Related Documentation

- Understanding JSRC-SAE Interactions on page 785
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 731
- Diameter Base Protocol Overview on page 703
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Gx-Plus for Provisioning Subscribers Overview on page 729

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 the master instance only in the context of the master routing instance

To configure Diameter base protocol:

1. Configure the origin realm and origin host of the Diameter master instance.

   See "Configuring the Origin Attributes of the Diameter Instance" on page 724

2. Configure the Diameter peers.

   See "Configuring Diameter Peers" on page 724

3. (Optional) Configure the Diameter transport layer elements.

   See "Configuring the Diameter Transport" on page 725

4. (Optional) Configure the Diameter network elements.

   See "Configuring Diameter Network Elements" on page 726

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

   See "Tracing Diameter Base Protocol Processes for Subscriber Access" on page 853.

Related Documentation

- Diameter Base Protocol Overview on page 703

Copyright © 2017, Juniper Networks, Inc.

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
   ```

Related Documentation

- Configuring Diameter on page 723
- origin on page 1234

Configuring Diameter Peers

You can configure the peers to which Diameter sends messages. By default, logical system *default* and routing instance *master* are used. 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# set peer p3
   ```

2. Specify the address of the Diameter peer.

   ```
   [edit diameter peer p3]
   user@host# set address 192.168.23.10
   ```

3. (Optional) Specify a routing instance, a logical system, or a logical system and routing instance for the Diameter peer.

   ```
   [edit diameter peer p3]
   user@host# set routing-instance ri8
   ```

   ```
   [edit diameter peer p3]
   user@host# set logical-system ls10
   ```

   ```
   [edit diameter peer p3]
   user@host# set logical-system ls10 routing-instance ri8
   ```


Copyright © 2017, Juniper Networks, Inc.

4. (Optional) Specify the port that Diameter uses for active connections to the peer.

   ```
   [edit diameter peer p3]
   user@host# set connect-actively port 49152
   ```

5. Specify the transport that Diameter uses for active connections to the peer.

   ```
   [edit diameter peer p3]
   user@host# set connect-actively transport t6
   ```

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 p3]
   user@host# set peer-origin host hostname
   user@host# set peer-origin realm realm name
   ```

Related Documentation

- Configuring Diameter on page 723

Configuring the Diameter Transport

You can configure one or more transports for a Diameter instance to set the IP address for the local connection, and optionally configure a logical system or routing instance context. By default, the logical system *default* and the routing instance *master* are used. 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# set transport t1
   ```

2. Configure the local IP address for the Diameter local transport connection.

   ```
   [edit diameter transport t1]
   user@host# set address 203.0.113.200
   ```

3. (Optional) Configure a logical system and optionally a routing instance for the transport.

   ```
   [edit diameter transport t1]
   user@host# set logical-system ls5
   ```

4. (Optional) Configure a routing instance for the transport.

   ```
   [edit diameter transport t1]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
user@host# set routing-instance ri10
```

Related Documentation

- Configuring Diameter on page 723

Configuring Diameter Network Elements

A Diameter network element (DNE) consists of associated functions, a list of prioritized peers, and a set of forwarding rules. The forwarding rules define individual routes through a set of associated destinations, functions, 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 724.

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 functions with the network element. All functions 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 a function (application) associated with the route.

Copyright © 2017, Juniper Networks, Inc.

```
[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>
```

Related Documentation

- Configuring Diameter on page 723

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 73

Configuring Gx-Plus for Provisioning Subscribers

- Gx-Plus for Provisioning Subscribers Overview on page 729
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 731
- Configuring Gx-Plus on page 736
- Configuring the Gx-Plus Partition on page 737
- Configuring Gx-Plus Global Attributes on page 738
- Provisioning Subscribers with Gx-Plus on page 739

Gx-Plus for Provisioning Subscribers Overview

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 60 on page 729 are interchangeable.

Table 60: 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

Copyright © 2017, Juniper Networks, Inc.

Table 60: Differences Between Gx-Plus and Junos OS Terminology *(continued)*

Gx-Plus	Junos OS
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

Copyright © 2017, Juniper Networks, Inc.

NOTE: More than one Diameter-based application (function), such as Gx-Plus or JSRC, can run on a router simultaneously.

Related Documentation

- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 731
- Configuring Gx-Plus on page 736

Understanding Gx-Plus Interactions Between the Router and the PCRF

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 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.

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)

Copyright © 2017, Juniper Networks, Inc.

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 fails unless the **authentication-order** statement in the profile specifies 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 services are activated, not even a default service.

Copyright © 2017, Juniper Networks, Inc.

- 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.

Copyright © 2017, Juniper Networks, Inc.

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 61 on page 734 lists router events and the subsequent router and PCRF actions.

Table 61: Router Events, Router Actions, and PCRF Actions

Router Event	Router Action	PCRF Action
The router receives no response from the PCRF or an error response.	Send event notification.	Respond to event notification.
The configuration changes. Significant changes such as the origin host or realm and the Gx-Plus partition destination host or realm also increment the value of the Origin-State-Id AVP.	Send event notification.	Respond to event notification and perform discovery.
The router receives an explicit discovery request from the PCRF.	Send event notification.	Respond to event notification.
The router undergoes a cold boot and all sessions are lost. This can result from a catastrophic failure or power cycle.	Send event notification.	Respond to event notification and clear the database.
The router undergoes a warm boot.	Send event notification.	Respond to event notification and clear the database.
Recovery resources that are needed to continuously retry unacknowledged requests (CCR-N and CCR-T messages) are exhausted. The value of the Origin-State-Id AVP is incremented. This event is unlikely to occur.	Send event notification.	Respond to event notification and perform discovery.

An important aspect of Gx-Plus fault tolerance is that subscriber login and termination requests are retried (replayed) forever until a satisfactory response is received from the PCRF. In rare circumstances, this can result in a stack of pending requests being replayed over and over.

You can issue the **clear network-access gx-plus replay** command to clear all pending requests. This command causes Gx-Plus to send a JSER message to PCRF that includes the Juniper-Event-Type AVP (AVP code 2103) with a value of 3 indicating a discovery request. The PCRF then returns a JDER message to initiate discovery of all subscribers. When this discovery completes, all pending subscriber requests are cleared.

Copyright © 2017, Juniper Networks, Inc.

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 service: Duration, input data, output data, or total data for the service 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 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 729
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Configuring Gx-Plus on page 736
- Default Subscriber Service Overview on page 247
- Configuring a Default Subscriber Service on page 248

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 737.

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 738.

3. Configure Gx-Plus provisioning for subscribers.

   See "Provisioning Subscribers with Gx-Plus" on page 739.

4. (Optional) Override PCRF control of a subscriber session to correct services or troubleshoot a problem.

Copyright © 2017, Juniper Networks, Inc.

See "Disabling PCRF Control of a Subscriber Session" on page 885.

5. (Optional) Configure Gx-Plus event tracing as part of general authentication service tracing operations.

   See "Tracing General Authentication Service Processes" on page 869.

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 729

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 723.

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]
   ```

Copyright © 2017, Juniper Networks, Inc.

```
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;
  }
}
```

Related Documentation

- Configuring Gx-Plus on page 736
- Gx-Plus for Provisioning Subscribers Overview on page 729

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
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Gx-Plus on page 736
- Gx-Plus for Provisioning Subscribers Overview on page 729

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. See "Configuring an Access Profile for Subscriber Management" on page 123.

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
```

Related Documentation

- Configuring Gx-Plus on page 736
- Gx-Plus for Provisioning Subscribers Overview on page 729

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 74

Configuring 3GPP Policy and Charging Control

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766
- Configuring the OCS Partition on page 769
- Configuring the PCRF Partition on page 773
- Configuring OCS Global Parameters on page 776

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

The 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) provides the unification of wireline provisioning and accounting for customers. Figure 23 on page 742 shows the components of an overall 3GPP PCC architecture.

Copyright © 2017, Juniper Networks, Inc.

Figure 23: 3GPP PCC Architecture Overview

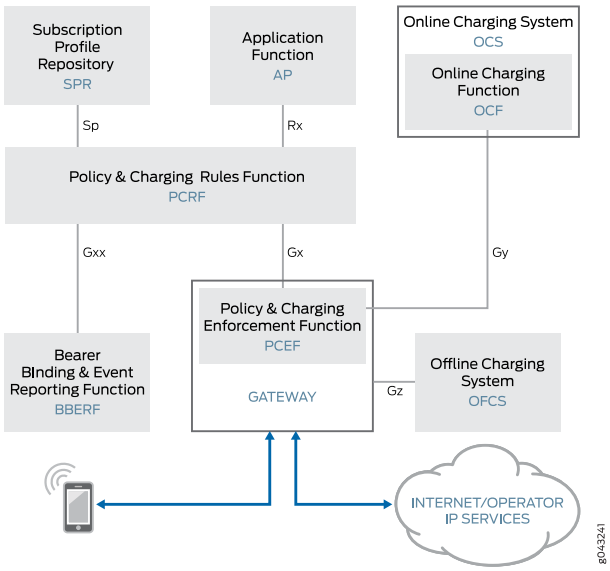

The four major components of the PCC architecture are:

- 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 62 on page 743 lists the functionality differences between PCRF and PCEF.

Copyright © 2017, Juniper Networks, Inc.

Table 62: 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 23 on page 742 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.

Related Documentation

- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

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 24 on page 744.

Figure 24: 3GPP PCC Architecture Overview

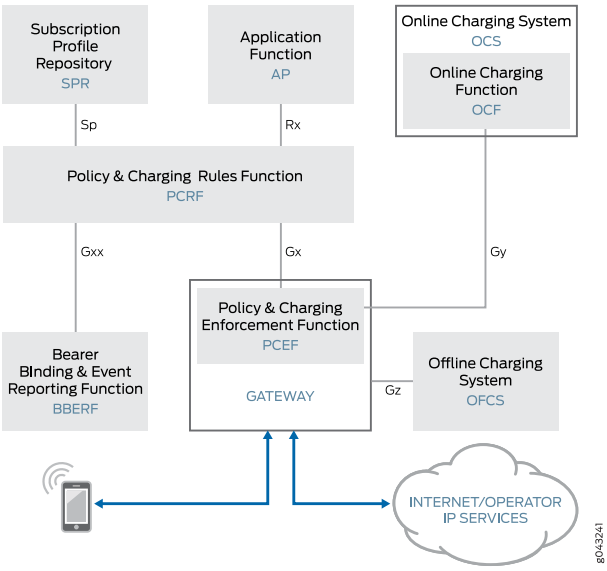

PCEF provides user traffic handling and quality of service (QoS) at the gateway, provides service data flow detection, and applies the rules received from the Policy Control and Charging Rules Function (PCRF). 3GPP defines Gx as the online policy protocol between the PCRF and the PCEF to provide control 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 on page 744
- Junos OS Environment on page 746

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.

Copyright © 2017, Juniper Networks, Inc.

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 63 on page 745 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 63: 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.

Copyright © 2017, Juniper Networks, Inc.

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 64 on page 746.

Table 64: 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

Copyright © 2017, Juniper Networks, Inc.

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-Defintion 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 65 on page 747 describes the initial set of fixed Juniper-Parameter VSAs.

Table 65: 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Understanding Gx Interactions Between the Router and the PCRF

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). They interact to perform the following subscriber access tasks:

- Subscriber Login on page 748
- Subscriber Update on page 752
- Subscriber Logout on page 753
- Subscriber Disconnect on page 754
- Connectivity Fault Recovery on page 755

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> }
}
}
```


Copyright © 2017, Juniper Networks, Inc.

```
[ 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 66 on page 749.

Table 66: Result-Code-AVP Categories

Result-Code-AVP Value	Result Category
SUCCESS(2001), LIMITED_SUCCSS(2002), and valid message	Grant

Copyright © 2017, Juniper Networks, Inc.

Table 66: Result-Code-AVP Categories *(continued)*

Result-Code-AVP Value	Result Category
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 67 on page 750, 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 67 on page 750 also contains PCRF local decision timeout expiration actions.

Table 67: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 67: 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 fails unless the **authentication-order** statement in the profile specifies 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.

Copyright © 2017, Juniper Networks, Inc.

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 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> }
```


Copyright © 2017, Juniper Networks, Inc.

```
[ 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 68 on page 753.

Table 68: 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
```

Copyright © 2017, Juniper Networks, Inc.

```
{ 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 69 on page 754.

Table 69: 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
```


Copyright © 2017, Juniper Networks, Inc.

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 70 on page 755.

Table 70: 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.

Copyright © 2017, Juniper Networks, Inc.

You can issue the **clear network-access pcrf subscribers** command to clear all PCRF subscribers.

Related Documentation

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Upstream and Downstream Messages for the PCRF on page 766
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Understanding Gy Interactions Between the Router and the OCS

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.

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 on page 756
- Intermediate Interrogation to the OCS on page 758
- Final Interrogation to the OCS on page 759
- Connectivity Fault Recovery on page 760
- Abort Session Requests on page 761

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:

Copyright © 2017, Juniper Networks, Inc.

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 }
    }
```

Copyright © 2017, Juniper Networks, Inc.

```
   { 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 }
```


Copyright © 2017, Juniper Networks, Inc.

```
 }
 { 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> }
```

Copyright © 2017, Juniper Networks, Inc.

```
[ 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 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Understanding Rule Interactions Between the PCRF, PCEF, and OCS

Figure 25 on page 761 shows the components of an overall 3rd Generation Partnership Project (3GPP) Policy and Charging Control (PCC) architecture.

Figure 25: 3GPP PCC Architecture Overview

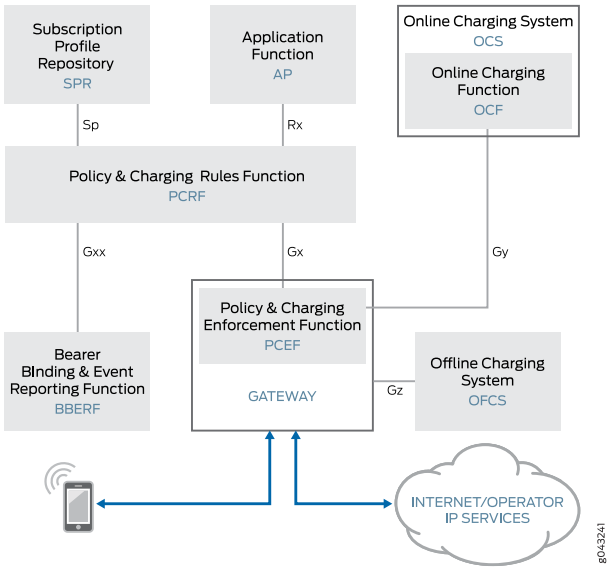

The Policy and Charging Rules Function (PCRF) pushes the rules down to the Policy and Charging Enforcement Function (PCEF) using the 3GPP Gx protocol and online policy interface. The PCEF provides service data flow detection, and applies the rules received from the PCRF. Optionally, the PCEF interacts with the Online Charging System (OCS) using the 3GPP Gy protocol to retrieve policy and charging authorization for quotas and

Copyright © 2017, Juniper Networks, Inc.

credit control. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

The behavior and rule interactions among the PCRF, PCEF, and OCS prohibit you from adding or modifying rules with a rating group, and require you to set the **force-continue** statement at the **[edit access ocs partition *partition-name*] hierarchy level**. Two loosely coupled finite state machines (FSMs), one for PCRF/Gx interactions and the other for OCS/Gy interactions, control the functionality.

The PCRF, PCEF, and OCS provide the following interactions:

- Login Interactions on page 762
- Update Interactions on page 762
- Quota Expiration and Validity-Time Interactions on page 763
- Connection and Monitoring Interactions on page 763
- Logout Interactions on page 765

Login Interactions

This login sequence of events is triggered by a CCA-GX-I message received between the PCRF and PCEF.

1. The PCRF installs the requested services.
2. The PCRF acknowledges subscriber activation and allows subscriber traffic to flow.
3. The PCRF sends a notification to the OCS.
4. The PCRF sends the first interrogation to the OCS using a CCR-GY-I message, and 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.
   - 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.
   - Otherwise, the system considers the response a no operation and takes no further action.

Update Interactions

This update sequence of events is triggered by an RAR-GX-U message received between the PCRF and PCEF.

Copyright © 2017, Juniper Networks, Inc.

1. If the update request contains any installation or modification of rules with rating groups, then the PCRF rejects the message request; otherwise, it acknowledges the message request.

2. The PCRF starts the service removal and installation process.

3. The PCRF waits for the service removal and installation process to complete, and if applicable, the OCS starts the final data collection process.
   - If the service being removed has a rating group, the service notifies the OCS that is part of the existing service removal process.
   - The OCS determines whether final data for this service is required to be reported, and then starts final data collection as necessary.
   - If data collection was started successfully, then the service moves into wait state.
4. The OCS waits for the final data collection process to complete. After final data collection completes or times out, the service notifies the OCS again. The OCS saves data in the subscriber's private directory and allows service to proceed with its existing service removal process.

5. 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.
   - 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.
   - Otherwise, the system considers the response a no operation and takes no further action.
6. The PCRF sends the intermediate interrogation to the OCS using a CCR-GY-U message.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

Logout Interactions

This logout sequence of events is triggered by a CCR-GX-T message to the PCRF that includes data for all unacknowledged rule-reports, and also sends a final interrogation to the OCS using a CCR-GY-T message that includes data for all unacknowledged final quota reports.

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, the service notifies the OCS that is part of the existing service removal process.
   - The OCS determines whether final data for this service is required to be reported, and then starts final data collection as necessary.
   - If data collection was started successfully, then the service moves into wait state.
3. The OCS waits for the final data collection process to complete. After final data collection completes or times out, the service notifies the OCS again. The OCS saves data in the subscriber's private directory and allows service to proceed with its existing service removal process.
4. The OCS allows the service deletion to proceed.
5. The OCS starts the final interrogation process.
6. Both the PCRF and the OCF waif for the final interrogation process to complete.
7. The OCF notifies the PCRF.
8. The PCRF starts a termination request, and then waits for the termination request to complete.
9. The PCRF allows the subscriber logout process to proceed.

Related Documentation

- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

Understanding Upstream and Downstream Messages for the PCRF

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. A subscriber can disconnect and reconnect, and it receives a different Session-Id.**

There is a one-to-one correspondence between a session and the Session-Id. 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 Code 263) is of type UTF8String.

```
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.

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 on page 767
- Common Downstream Messages on page 768

 Copyright © 2017, Juniper Networks, Inc.

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 on page 767
- Charging Rules Installation Notifications on page 767
- Event Trigger Commands on page 768

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 71 on page 767 varies depending on the transaction.

Table 71: 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) }
  }
```

Copyright © 2017, Juniper Networks, Inc.

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 on page 768
- Charging Rule Removal Commands on page 768
- Event Trigger Commands on page 769

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” }
```


Copyright © 2017, Juniper Networks, Inc.

```
  { 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.

Related Documentation

- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744

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 723.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```


Copyright © 2017, Juniper Networks, Inc.

10. Configure the OCS partition so that subscriber traffic is allowed before the first OCS interrogation and services are not removed by the PCEF when it receives negative responses from the OCS.

    [edit access ocs partition *partition-name*]
    user@host# set force-continue

11. (Optional) Configure the GGSN-Address AVP value used in all CCR-GY messages.

    [edit access ocs partition *partition-name*]
    user@host# set ggsn-address ***address***

12. (Optional) Configure the 3GPP-GGSN-MCC-MNC AVP value used in all CCR-GY messages.

    [edit access ocs partition *partition-name*]
    user@host# set ggsn-mcc-mnc***ggsn-mcc-mnc***

13. (Optional) Configure the number of outstanding requests from the OCS to the OCS server that can be retried when the requests are improperly answered.

    [edit access ocs partition *partition-name*]
    user@host# set max-outstanding-requests ***number***

14. (Optional) Specify that the Origin-State-ID AVP is included in Diameter base protocol level messages for the partition, and synchronized with the latest value sent to aid in monitoring changes in value.

    [edit access ocs partition *partition-name*]
    user@host# set send-origin-state-id

15. (Optional) Configure the information concatenated as a string in usernames that the OCS partition sends to the PCEF to identify the subscribers.

    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*]

Copyright © 2017, Juniper Networks, Inc.

```
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
```

Related Documentation

- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

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 723.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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. Configure the PCRF partition so that the subscriber login proceeds even if the PCRF is unavailable or the PCRF does not respond within the corresponding logout response timeout period.

```
[edit access pcrf partition partition-name]
user@host# set local-decision
```

9. (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
```

10. (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
```

11. (Optional) Specify that the PCRF sends local report downstream messages by default.

```
[edit access pcrf partition partition-name]
user@host# set report-local-rule
```

12. (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
```

13. (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
```

14. (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
```


Copyright © 2017, Juniper Networks, Inc.

15. (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

16. (Optional) Specify that the

    [edit access pcrf partition *partition-name*]
    user@host# **set** send-origin-state-id

17. (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

18. (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**

Copyright © 2017, Juniper Networks, Inc.

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
```

19. (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
```

20. (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
```

Related Documentation

- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

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]
```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# set service-context-id service-context
```

Related Documentation

- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 75

Configuring NASREQ for Authentication and Authorization

- Diameter Network Access Server Application (NASREQ) on page 779
- Configuring the Diameter Network Access Server Application (NASREQ) on page 781

Diameter Network Access Server Application (NASREQ)

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:

Copyright © 2017, Juniper Networks, Inc.

- *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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Diameter Base Protocol Overview on page 703
- Configuring the Diameter Network Access Server Application (NASREQ) on page 781
- Diameter AVPs and Diameter Applications on page 711

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*

Copyright © 2017, Juniper Networks, Inc.

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 Network Access Server Application (NASREQ) on page 779
- Diameter Base Protocol Overview on page 703
- Diameter AVPs and Diameter Applications on page 711

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 76

Configuring JSRC in Subscriber Access Networks

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Understanding JSRC-SAE Interactions on page 785
- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788
- Assigning a Partition to JSRC on page 789
- Authorizing Subscribers with JSRC on page 789
- Provisioning Subscribers with JSRC on page 790

Juniper Networks Session and Resource Control (SRC) and JSRC Overview

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Hardware Requirements for JSRC for Subscriber Access

JSRC is supported on Juniper Networks MX Series 3D Universal Edge Routers. JSRC currently supports subscriber sessions on static and dynamic interfaces.

Related Documentation

- Understanding JSRC-SAE Interactions on page 785
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Configuring JSRC on page 787

Copyright © 2017, Juniper Networks, Inc.

Understanding JSRC-SAE Interactions

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 is enabled for DHCP subscribers when you include the **authorization-order jsrc** statement at the **[edit access profile *profile-name*]** hierarchy level. This setting causes AAA to ignore the authentication order setting in the access profile. As a result, AAA does not authenticate the DHCP subscribers. For non-DHCP subscribers, AAA ignores the **authorization-order** statement.

When a DHCP subscriber attempts to log in, DHCP 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711
- Configuring JSRC on page 787

Configuring JSRC

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 788.

2. Assign the JSRC partition.

   See "Assigning a Partition to JSRC" on page 789.

3. Configure JSRC authorization for subscribers.

   See "Authorizing Subscribers with JSRC" on page 789.

4. Configure JSRC provisioning for subscribers.

   See "Provisioning Subscribers with JSRC" on page 790.

5. (Optional) Configure JSRC to exclude an AVP from Messages Sent to SAE.

   See "Excluding AVPs from Diameter Messages for JSRC" on page 791.

Copyright © 2017, Juniper Networks, Inc.

6. Configure service accounting by JSRC.

   See "Configuring Service Accounting with JSRC" on page 794.

7. Configure JSRC event tracing as part of general authentication service tracing operations.

   See "Tracing General Authentication Service Processes" on page 869.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783

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 723.

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]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@host# set destination-realm generic.example.com
```

Related Documentation

- Configuring JSRC on page 787

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 788

To assign the JSRC partition:

- Specify the partition name.

```
[edit jsrc]
user@host# set jsrc-partition partition1
```

Related Documentation

- Configuring JSRC on page 787

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
```

Related Documentation

- Configuring JSRC on page 787
- Creating Unique Usernames for DHCP Clients on page 282
- profile on page 1278

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- Configuring JSRC on page 787

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 77

Excluding Diameter AVPs from JSRC Messages

- Excluding AVPs from Diameter Messages for JSRC on page 791

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
   ```

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
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

- Understanding JSRC-SAE Interactions on page 785

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 78

Configuring Service Accounting with JSRC

- Service Accounting with JSRC on page 793
- Configuring Service Accounting with JSRC on page 794

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Service Packet Counting on page 109
- Messages Used by Diameter Applications on page 706
- Diameter AVPs and Diameter Applications on page 711

Configuring Service Accounting with JSRC

You can configure JSRC to report accounting statistics for service sessions.

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.

To configure service accounting by JSRC:

1. Configure JSRC to provision subscriber services.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

Related Documentation

- Service Accounting with JSRC on page 793

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 79

Configuring Subscribers on Static Interfaces

- Subscribers on Static Interfaces Overview on page 797
- Configuring Subscribers over Static Interfaces on page 800
- Example: Configuring Static Subscribers for Subscriber Access on page 801

Subscribers on Static Interfaces Overview

You can associate subscribers with statically configured interfaces and provide dynamic service activation and activation 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. 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.

Alternatively, 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.

Currently, only Ethernet interfaces support static subscribers. 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**

Copyright © 2017, Juniper Networks, Inc.

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, 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 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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- Understanding JSRC-SAE Interactions on page 785

Configuring Subscribers over Static Interfaces

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 805.

2. Specify the global dynamic profile that is instantiated when static subscribers log in.

   See "Specifying the Static Subscriber Global Dynamic Profile" on page 805.

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 806

4. Configure the global authentication password for static subscribers.

   See "Configuring the Static Subscriber Global Authentication Password" on page 807.

5. Configure the global username for static subscribers.

   See "Configuring the Static Subscriber Global Username" on page 807.

6. Configure a group of subscribers to share values different from the global configuration.

   See "Creating a Static Subscriber Group" on page 809.

7. Specify the access profile for the static subscriber group.

Copyright © 2017, Juniper Networks, Inc.

See "Specifying the Static Subscriber Group Access Profile" on page 810.

8. Specify the dynamic profile for the static subscriber group.

   See "Specifying the Static Subscriber Group Dynamic Profile" on page 810.

9. 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 811.

10. Configure the authentication password for the static subscriber group.

    See "Configuring the Static Subscriber Group Authentication Password" on page 811.

11. Configure the username for the static subscriber group.

    See "Configuring the Static Subscriber Group Username" on page 812.

12. (Optional) Force a static subscriber to be logged out from an interface.

    See "Forcing a Static Subscriber to Be Logged Out" on page 819.

13. (Optional) Enable an interface to accept static subscriber logins.

    See "Resetting the State of an Interface for Static Subscriber Login" on page 819.

14. (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 820.

15. (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 820.

16. Configure trace options for troubleshooting the configuration.

    See "Tracing Static Subscriber Operations" on page 881.

Related Documentation

- Subscribers on Static Interfaces Overview on page 797

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;
      }
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
}
```

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;
  }
}
```

Related Documentation

- Subscribers on Static Interfaces Overview on page 797
- Configuring Subscribers over Static Interfaces on page 800

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 80

Configuring the Static Subscribers Global Profile

- Specifying the Static Subscriber Global Access Profile on page 805
- Specifying the Static Subscriber Global Dynamic Profile on page 805
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers on page 806
- Configuring the Static Subscriber Global Authentication Password on page 807
- Configuring the Static Subscriber Global Username on page 807

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
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Group Access Profile on page 810
- profile on page 1278

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.

Copyright © 2017, Juniper Networks, Inc.

```
[edit system services static-subscribers]
user@host# set dynamic-profile dyn-profile-1
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Group Dynamic Profile on page 810
- *dynamic-profiles*

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.

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
  ```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Group Dynamic Profile on page 810
- dynamic-profile on page 1070

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Group Authentication Password on page 811
- authentication on page 958

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 five 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, and routing instance name are derived from the configuration context. The elements are ordered as follows:

user-prefix.interface.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 Building5
   ```

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 logical system name is included in the username.

   ```
   [edit system services static-subscribers authentication username-include]
   user@host# set logical-system-name
   ```

4. Specify that the routing instance name is included in the username.

Copyright © 2017, Juniper Networks, Inc.

```
[edit system services static-subscribers authentication username-include]
user@host# set routing-instance-name
```

5. Specify the domain name included in the username.

```
[edit system services static-subscribers authentication username-include]
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

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Group Username on page 812
- username-include on page 1458

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 81

Configuring the Static Subscribers Group Profile

- Creating a Static Subscriber Group on page 809
- Specifying the Static Subscriber Group Access Profile on page 810
- Specifying the Static Subscriber Group Dynamic Profile on page 810
- Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group on page 811
- Configuring the Static Subscriber Group Authentication Password on page 811
- Configuring the Static Subscriber Group Username on page 812

Creating a Static Subscriber Group

You can override the configuration that is applied globally to static subscribers by creating a static subscriber group that consists of a set of statically configured interfaces. You can then apply a common configuration for the group with values different from the global values for access and dynamic profiles, password, and username.

To configure an interface group for static subscribers:

1. Access the **[edit system services static-subscribers]** hierarchy level.

2. Create the group and assign the name.

   ```
   [edit system services static-subscribers]
   user@host# edit group boston
   ```

3. Specify the names of one or more interfaces on which static subscribers can be created. You can repeat the interface ***interface-name*** statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

   ```
   [edit system services static-subscribers group boston]
   user@host# set interface ge-1/0/1.1
   user@host# set interface ge-1/0/1.2
   ```

4. (Optional) You can use the **upto** ***upto-interface-name*** option to specify a range of interfaces for a group.

Copyright © 2017, Juniper Networks, Inc.

```
[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
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Group Access Profile on page 810
- Specifying the Static Subscriber Group Dynamic Profile on page 810
- Configuring the Static Subscriber Group Authentication Password on page 811
- Configuring the Static Subscriber Group Username on page 812

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
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- profile on page 1278

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
```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Global Dynamic Profile on page 805
- *dynamic-profiles*

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
  ```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Group Dynamic Profile on page 810
- dynamic-profile on page 1070

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.

Copyright © 2017, Juniper Networks, Inc.

```
[edit system services static-subscribers group boston authentication]
user@host# set password $ABC123
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Authentication Password on page 807
- authentication on page 958

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 five 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, and routing instance name are derived from the configuration context. The elements are ordered as follows:

user-prefix*.*interface*.*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 boston authentication username-include]
   user@host# set user-prefix 2ndFloor
   ```

2. (Optional) Specify that the interface name is included in the username.

   ```
   [edit system services static-subscribers group boston authentication username-include]
   user@host# set interface
   ```

3. (Optional) Specify that the logical system name is included in the username.

   ```
   [edit system services static-subscribers group boston authentication username-include]
   user@host# set logical-system-name
   ```

4. Specify that the routing instance name is included in the username.

   ```
   [edit system services static-subscribers group boston authentication username-include]
   user@host# set routing-instance-name
   ```

5. Specify the domain name included in the username.

   ```
   [edit system services static-subscribers group boston authentication username-include]
   user@host# set domain-name building5.example.com
   ```

Configured in the default logical system and master routing instance for interface ge-0/1/2.50, this sample configuration generates the following username:

Copyright © 2017, Juniper Networks, Inc.

2ndfloor.ge-0-1-2-50.default.master.building5.example.com

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- username-include on page 1458

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 82

Monitoring and Managing Diameter Information for Subscriber Access

- Verifying Diameter Node, Instance, and Route Information on page 815
- Verifying and Managing Diameter Function Information on page 816
- Verifying and Managing Diameter Peer Information on page 817
- Verifying Diameter Network Element Information on page 818

Verifying Diameter Node, Instance, and Route Information

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 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**

Copyright © 2017, Juniper Networks, Inc.

- 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**

Related Documentation

- Configuring Diameter on page 723
- Configuring Gx-Plus on page 736
- CLI Explorer

Verifying and Managing Diameter Function Information

Purpose View or clear Diameter function information:

Action

- To display summary information about all functions associated with Diameter:

  user@host> **show diameter function**

- To display detailed information about all functions associated with Diameter:

  user@host> **show diameter function detail**

- To display information about a specific function associated with Diameter, add the function name to the command:

  user@host> **show diameter function jsrc**

  user@host> **show diameter function detail gx-plus**

- To display summary statistics about all functions associated with Diameter:

  user@host> **show diameter function statistics**

- To display detailed statistics about all functions associated with Diameter:

  user@host> **show diameter function statistics detail**

- To display statistics about a specific function associated with Diameter, add the function 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 functions associated with Diameter:

  user@host>**clear diameter function statistics**

- To delete current statistics for a specific function associated with Diameter:

  user@host>**clear diameter function gx-plus statistics**

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Diameter on page 723
- Configuring Gx-Plus on page 736
- CLI Explorer

Verifying and Managing Diameter Peer Information

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**

Copyright © 2017, Juniper Networks, Inc.

- To delete the specified Diameter peer and its current statistics:

  user@host>**clear diameter peer peer5 statistics**

Related Documentation

- Configuring Diameter on page 723
- CLI Explorer

Verifying Diameter Network Element Information

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

- Configuring Diameter on page 723
- CLI Explorer

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 83

Monitoring and Managing Subscriber Information on Static Interfaces

- Forcing a Static Subscriber to Be Logged Out on page 819
- Resetting the State of an Interface for Static Subscriber Login on page 819
- Forcing a Group of Static Subscribers to Be Logged Out on page 820
- Resetting the State of an Interface Group for Static Subscriber Login on page 820
- Verifying Information about Subscriber Sessions on Static Interfaces on page 820

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**

Related Documentation

- Resetting the State of an Interface for Static Subscriber Login on page 819

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**

Related Documentation

- Forcing a Static Subscriber to Be Logged Out on page 819

Copyright © 2017, Juniper Networks, Inc.

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**

Related Documentation

- Resetting the State of an Interface Group for Static Subscriber Login on page 820

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**

Related Documentation

- Forcing a Group of Static Subscribers to Be Logged Out on page 820

Verifying Information about Subscriber Sessions on Static Interfaces

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

- For more information, see the CLI Explorer
- Configuring Subscribers over Static Interfaces on page 800
- Subscribers on Static Interfaces Overview on page 797

Copyright © 2017, Juniper Networks, Inc.

PART 10

Troubleshooting

- Configuring AAA Testing and Troubleshooting on page 823
- Tracing Extended DHCP Operations on page 829
- Configuring Subscriber Management Database Log Files on page 837
- Configuring Subscriber Management Database Trace Flags and Operations on page 839
- Configuring Subscriber Management Session Database Log Files on page 843
- Configuring Subscriber Management Session Database Trace Flags and Operations on page 845
- Configuring Diameter Base Protocol Log Files on page 849
- Configuring Diameter Base Protocol Trace Flags and Operations on page 853
- Troubleshooting Diameter Networks on page 857
- Configuring ANCP Log Files on page 859
- Configuring ANCP Trace Flags and Operations on page 863
- Configuring General Authentication Service Log Files on page 867
- Configuring General Authentication Service Trace Flags and Operations on page 869
- Configuring Static Subscriber Interfaces Log Files on page 877
- Configuring Static Subscriber Interfaces Trace Flags and Operations on page 881
- Overriding PCRF Session Control to Troubleshoot a Session or Services on page 885
- Contacting Juniper Technical Support on page 887

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 84

Configuring AAA Testing and Troubleshooting

- AAA Configuration Testing and Troubleshooting on page 823
- Testing a Subscriber AAA Configuration on page 823

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 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.

Related Documentation

- Testing a Subscriber AAA Configuration on page 823

Testing a Subscriber AAA Configuration

Purpose Display the AAA attributes that subscriber management assigns to the subscriber during login.

Copyright © 2017, Juniper Networks, Inc.

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.

NOTE: 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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

```
        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
```


Copyright © 2017, Juniper Networks, Inc.

```
    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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
     ...
```

Related Documentation

- AAA Configuration Testing and Troubleshooting on page 823

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 85

Tracing Extended DHCP Operations

- Tracing Extended DHCP Operations on page 829

Tracing Extended DHCP Operations

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.

Copyright © 2017, Juniper Networks, Inc.

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 834.

The extended DHCP traceoptions operations are described in the following sections:

- Configuring the Extended DHCP Log Filename on page 831
- Configuring the Number and Size of Extended DHCP Log Files on page 831
- Configuring Access to the Extended DHCP Log File on page 832
- Configuring a Regular Expression for Extended DHCP Messages to Be Logged on page 832
- Configuring the Extended DHCP Tracing Flags on page 832
- Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged on page 833
- Tracing Extended DHCP Operations for Specific Interfaces on page 834

Copyright © 2017, Juniper Networks, Inc.

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*

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Copyright © 2017, Juniper Networks, Inc.

```
[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 829.

   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 831.

      - Configure the number and size of the log files.

        See "Configuring the Number and Size of Extended DHCP Log Files" on page 831.

      - Configure access to the log file.

        See "Configuring Access to the Extended DHCP Log File" on page 832.

      - Configure a regular expression to filter logging events.

        See "Configuring a Regular Expression for Extended DHCP Messages to Be Logged" on page 832.

   c. (Optional) Specify tracing flag options.

      See "Configuring the Extended DHCP Tracing Flags" on page 832.

   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 833.

Copyright © 2017, Juniper Networks, Inc.

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
     ```

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 86

Configuring Subscriber Management Database Log Files

- Configuring the Number and Size of Subscriber Management Database Log Files on page 837
- Configuring Access to the Subscriber Management Database Log File on page 838
- Configuring a Regular Expression for Subscriber Management Database Messages to Be Logged on page 838

Configuring the Number and Size of Subscriber Management Database 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 services subscriber-management traceoptions]
  user@host# set file smi_1 _logfile_1 files 20 size 2097152

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Copyright © 2017, Juniper Networks, Inc.

Configuring Access to the Subscriber Management Database 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 services subscriber-management traceoptions]
  user@host# **set file smi_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 services subscriber-management traceoptions]
  user@host# **set file smi_1 _logfile_1 no-world-readable**

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Configuring a Regular Expression for Subscriber Management Database 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 services subscriber-management traceoptions]
  user@host# **set file smi_1 _logfile_1 match *regex***

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 87

Configuring Subscriber Management Database Trace Flags and Operations

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839
- Configuring the Subscriber Management Database Trace Log Filename on page 840
- Configuring the Subscriber Management Database Tracing Flags on page 840

Tracing Subscriber Management Database Operations for Subscriber Access

The Junos OS trace feature tracks subscriber management database 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 **smid**. 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.

Copyright © 2017, Juniper Networks, Inc.

To configure all aspects of subscriber management database tracing operations:

1. Configure a trace log filename.

   See "Configuring the Subscriber Management Database Trace Log Filename" on page 840.

2. Configure the number and size of trace logs.

   See "Configuring the Number and Size of Subscriber Management Database Log Files" on page 837.

3. Configure user access to trace logs.

   See "Configuring Access to the Subscriber Management Database Log File" on page 838.

4. Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Subscriber Management Database Messages to Be Logged" on page 838.

5. Configure flags to specify which events are logged.

   See "Configuring the Subscriber Management Database Tracing Flags" on page 840.

Configuring the Subscriber Management Database Trace Log Filename

By default, the name of the file that records trace output for the subscriber management database is **smid**. You can specify a different name with the **file** option.

To configure the filename for subscriber management database tracing operations:

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

  ```
  [edit system services subscriber-management traceoptions]
  user@host# set file smi_logfile_1
  ```

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Configuring the Subscriber Management Database 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
  ```


Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 88

Configuring Subscriber Management Session Database Log Files

- Configuring the Number and Size of Subscriber Management Session Database Replication Log Files on page 843
- Configuring Access to the Subscriber Management Session Database Replication Log File on page 844
- Configuring a Regular Expression for Subscriber Management Session Database Replication Messages to Be Logged on page 844

Configuring the Number and Size of Subscriber Management Session Database Replication 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 services database-replication traceoptions]
  user@host# **set file bdrep_1 _logfile_1 files 20 size 2097152**

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Configuring Access to the Subscriber Management Session Database Replication 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 services database-replication traceoptions]
user@host# set file bdrep_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 services database-replication traceoptions]
user@host# set file bdrep_1 _logfile_1 no-world-readable
```

Related Documentation
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Configuring a Regular Expression for Subscriber Management Session Database Replication 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 services database-replication traceoptions]
user@host# set file bdrep_1 _logfile_1 match regex
```

Related Documentation
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 89

Configuring Subscriber Management Session Database Trace Flags and Operations

- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845
- Configuring the Subscriber Management Session Database Replication Trace Log Filename on page 846
- Configuring the Subscriber Management Session Database Replication Tracing Flags on page 846

Tracing Subscriber Management Session Database Replication Operations for Subscriber Access

The Junos OS trace feature tracks subscriber management database replication 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 **bdbrepd**. 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 configure the maximum file size to be from 10 KB through 1 gigabyte (GB). You can also specify the number of trace files to be from 2 through 1000. (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.

Copyright © 2017, Juniper Networks, Inc.

To configure all aspects of subscriber management database replication tracing operations:

1. Configure a trace log filename.

   See "Configuring the Subscriber Management Session Database Replication Trace Log Filename" on page 846.

2. Configure the number and size of trace logs.

   See "Configuring the Number and Size of Subscriber Management Session Database Replication Log Files" on page 843.

3. Configure user access to trace logs.

   See "Configuring Access to the Subscriber Management Session Database Replication Log File" on page 844.

4. Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Subscriber Management Session Database Replication Messages to Be Logged" on page 844.

5. Configure flags to specify which events are logged.

   See "Configuring the Subscriber Management Session Database Replication Tracing Flags" on page 846.

Configuring the Subscriber Management Session Database Replication Trace Log Filename

By default, the name of the file that records trace output for the subscriber management database is **bdbrepd**. You can specify a different name with the **file** option.

To configure the filename for subscriber management database tracing operations:

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

  ```
  [edit system services database-replication traceoptions]
  user@host# set file bdbrep_logfile_1
  ```

Related Documentation

- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Configuring the Subscriber Management Session Database Replication 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.

Copyright © 2017, Juniper Networks, Inc.

To configure the flags for the events to be logged:

- Configure the flags.

  ```
  [edit system services database-replication traceoptions]
  user@host# set flag flag
  ```

Related Documentation

- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 90

Configuring Diameter Base Protocol Log Files

- Configuring the Number and Size of Diameter Base Protocol Log Files on page 849
- Configuring Access to the Diameter Base Protocol Log File on page 850
- Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged on page 850
- Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged on page 850

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
  ```

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

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
```

Related Documentation
- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

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
```

Related Documentation
- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

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

Copyright © 2017, Juniper Networks, Inc.

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

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 91

Configuring Diameter Base Protocol Trace Flags and Operations

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853
- Configuring the Diameter Base Protocol Trace Log Filename on page 854
- Configuring the Diameter Base Protocol Tracing Flags on page 854

Tracing Diameter Base Protocol Processes for Subscriber Access

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.

Copyright © 2017, Juniper Networks, Inc.

To configure Diameter base protocol tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Diameter Base Protocol Trace Log Filename" on page 854.

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

   See "Configuring the Number and Size of Diameter Base Protocol Log Files" on page 849.

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

   See "Configuring Access to the Diameter Base Protocol Log File" on page 850.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged" on page 850.

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

   See "Configuring the Diameter Base Protocol Tracing Flags" on page 854.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged" on page 850.

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
  ```

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

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.

Copyright © 2017, Juniper Networks, Inc.

To configure the flags for the events to be logged:

- Configure the flags.

  ```
  [edit system processes diameter-service traceoptions]
  user@host# set flag dne
  ```

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 92

Troubleshooting Diameter Networks

- Troubleshooting Diameter Network Configuration on page 857
- Troubleshooting Diameter Network Connectivity on page 858

Troubleshooting Diameter Network Configuration

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.

Related Documentation

- show diameter function statistics on page 1676
- show diameter network-element map on page 1684
- show diameter peer on page 1687
- show diameter peer map on page 1692
- show diameter route on page 1699

Copyright © 2017, Juniper Networks, Inc.

Troubleshooting Diameter Network Connectivity

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

- show diameter peer on page 1687
- show diameter peer statistics on page 1695

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 93

Configuring ANCP Log Files

- Configuring the Number and Size of ANCP Log Files on page 859
- Configuring Access to the ANCP Log File on page 860
- Configuring a Regular Expression for ANCP Messages to Be Logged on page 860
- Configuring the Severity Level to Filter Which ANCP Messages Are Logged on page 860

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**

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

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
  ```

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

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**.

Copyright © 2017, Juniper Networks, Inc.

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

- Tracing ANCP Agent Operations for Subscriber Access on page 863
- Configuring the ANCP Agent on page 648
- traceoptions on page 1410

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 94

Configuring ANCP Trace Flags and Operations

- Tracing ANCP Agent Operations for Subscriber Access on page 863
- Configuring the ANCP Trace Log Filename on page 864
- Configuring the ANCP Tracing Flags on page 864

Tracing ANCP Agent Operations for Subscriber Access

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.

Copyright © 2017, Juniper Networks, Inc.

To configure ANCP agent tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the ANCP Trace Log Filename" on page 864.

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

   See "Configuring the Number and Size of ANCP Log Files" on page 859.

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

   See "Configuring Access to the ANCP Log File" on page 860.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for ANCP Messages to Be Logged" on page 860.

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

   See "Configuring the ANCP Tracing Flags" on page 864.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which ANCP Messages Are Logged" on page 860.

Related Documentation

- Configuring the ANCP Agent on page 648

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
```

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

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.

Copyright © 2017, Juniper Networks, Inc.

To configure the flags for the events to be logged:

- Configure the flags.

  ```
  [edit protocols ancp traceoptions]
  user@host# set flag restart
  ```

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 95

Configuring General Authentication Service Log Files

- Configuring the Number and Size of General Authentication Service Processes Log Files on page 867
- Configuring Access to the Log File on page 868
- Configuring a Regular Expression for Lines to Be Logged on page 868

Configuring the Number and Size of General Authentication Service Processes 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
  ```

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Tracing General Authentication Service Processes on page 869

Configuring Access to the 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
```

Related Documentation
- Tracing General Authentication Service Processes on page 869

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events. You can refine the output by including 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
```

Related Documentation
- Tracing General Authentication Service Processes on page 869

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 96

Configuring General Authentication Service Trace Flags and Operations

- Tracing General Authentication Service Processes on page 869
- Configuring the General Authentication Service Processes Trace Log Filename on page 873
- Configuring the Trace Operation on page 874
- Configuring Subscriber Filtering for General Authentication Service Trace Operations on page 874

Tracing General Authentication Service Processes

The Junos OS trace operations feature tracks general authentication service operations and records events in a log file. By default, the tracing operation is inactive. To trace general authentication service processes, you specify flags in the **traceoptions** statement at the **[edit system processes general-authentication-service]** hierarchy level. The default tracing behavior is the following:

- 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.
- 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.

Copyright © 2017, Juniper Networks, Inc.

- 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 general authentication service tracing operations are described in the following sections:

- Configuring the General Authentication Service Processes Trace Log Filename on page 870
- Configuring the Number and Size of General Authentication Service Processes Log Files on page 870
- Configuring Access to the Log File on page 871
- Configuring a Regular Expression for Lines to Be Logged on page 871
- Configuring Subscriber Filtering for General Authentication Service Trace Operations on page 871
- Configuring the Trace Operation on page 872

Configuring the General Authentication Service Processes 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 Processes 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**.

Copyright © 2017, Juniper Networks, Inc.

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 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 Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events. You can refine the output by including 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 Trace Operations

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.

Copyright © 2017, Juniper Networks, Inc.

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 Trace Operation

By default, only important events are logged. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
address-assignment	Trace all address-assignment pool events

Copyright © 2017, Juniper Networks, Inc.

Flag	Description
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

To configure the flags for the event to be logged:

- Configure the flags.

  [edit system processes general-authentication-service traceoptions]
  user@host# set flag address-assignment

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.

Configuring the General Authentication Service Processes 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

Related Documentation

- Tracing General Authentication Service Processes on page 869
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

Configuring the Trace Operation

By default, only important events are logged. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
address-assignment	Trace all 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

To configure the flags for the event to be logged:

- Configure the flags.

  [edit system processes general-authentication-service traceoptions]
  user@host# set flag address-assignment

Related Documentation

- Tracing General Authentication Service Processes on page 869

Configuring Subscriber Filtering for General Authentication Service Trace Operations

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

Copyright © 2017, Juniper Networks, Inc.

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
  ```

Copyright © 2017, Juniper Networks, Inc.

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

- Tracing General Authentication Service Processes on page 869

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 97

Configuring Static Subscriber Interfaces Log Files

- Configuring the Number and Size of Static Subscribers Log Files on page 877
- Configuring Access to the Static Subscribers Log File on page 878
- Configuring a Regular Expression for Static Subscriber Messages to Be Logged on page 878
- Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged on page 878

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

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Tracing Static Subscriber Operations on page 881

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
```

Related Documentation
- Tracing Static Subscriber Operations on page 881

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
```

Related Documentation
- Tracing Static Subscriber Operations on page 881

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

Copyright © 2017, Juniper Networks, Inc.

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

- Tracing Static Subscriber Operations on page 881

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 98

Configuring Static Subscriber Interfaces Trace Flags and Operations

- Tracing Static Subscriber Operations on page 881
- Configuring the Static Subscribers Trace Log Filename on page 882
- Configuring the Static Subscribers Tracing Flags on page 882

Tracing Static Subscriber Operations

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.

Copyright © 2017, Juniper Networks, Inc.

To configure static subscriber tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Static Subscribers Trace Log Filename" on page 882.

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

   See "Configuring the Number and Size of Static Subscribers Log Files" on page 877.

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

   See a"Configuring Access to the Static Subscribers Log File" on page 878.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Static Subscriber Messages to Be Logged" on page 878.

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

   See "Configuring the Static Subscribers Tracing Flags" on page 882.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged" on page 878.

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
  ```

Related Documentation

- Tracing Static Subscriber Operations on page 881

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.

Copyright © 2017, Juniper Networks, Inc.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit system processes static-subscribers traceoptions]
  user@host# **set flag authentication**

Related Documentation

- Tracing Static Subscriber Operations on page 881

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 99

Overriding PCRF Session Control to Troubleshoot a Session or Services

- Disabling PCRF Control of a Subscriber Session on page 885

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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Related Documentation

- Activating and Deactivating Subscriber Services Locally with the CLI on page 610

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 100

Contacting Juniper Technical Support

- Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support on page 887
- Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support on page 889

Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support

Problem **Description:** When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper Networks Technical Support. This topic shows you the most useful logs for a variety of network implementations. In addition to the relevant log information, you must also collect standard troubleshooting information and send it to Juniper Networks Technical Support in your request for assistance.

Solution To collect standard troubleshooting information:

- Redirect the command output to a file.

  user@host> **request support information | save rsi-1**

Copyright © 2017, Juniper Networks, Inc.

To configure logging to assist Juniper Networks Technical Support:

1. Review the following blocks of statements to determine which apply to your configuration.

```
[edit]
set system syslog archive size 100m files 25
set system auto-configuration traceoptions file filename
set system auto-configuration traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions level all
set protocols ppp-service traceoptions flag all
set protocols ppp traceoptions file filename size 100m files 25
set protocols ppp traceoptions level all
set protocols ppp traceoptions flag all
set protocols ppp monitor-session all
set interfaces pp0 traceoptions flag all
set demux traceoptions file filename size 100m files 25
set demux traceoptions level all
set demux traceoptions flag all
set system processes dhcp-service traceoptions file filename
set system processes dhcp-service traceoptions file size 100m
set system processes dhcp-service traceoptions file files 25
set system processes dhcp-service traceoptions flag all
set class-of-service traceoptions file filename
set class-of-service traceoptions file size 100m
set class-of-service traceoptions flag all
set class-of-service traceoptions file files 25
set routing-options traceoptions file filename
set routing-options traceoptions file size 100m
set routing-options traceoptions flag all
set routing-options traceoptions file files 25
set interfaces traceoptions file filename
set interfaces traceoptions file size 100m
set interfaces traceoptions flag all
set interfaces traceoptions file files 25
set system processes general-authentication-service traceoptions file filename
set system processes general-authentication-service traceoptions file size 100m
set system processes general-authentication-service traceoptions flag all
set system processes general-authentication-service traceoptions file files 25
```

2. Copy the relevant statements into a text file and modify the log filenames as you want.

3. Copy the statements from the text file and paste them into the CLI on your router to configure logging.

4. Commit the logging configuration to begin collecting information.

NOTE: The maximum file size for DHCP local server and DHCP relay log files is 1 GB. The maximum number of log files for DHCP local server and DHCP relay is 1000.

Copyright © 2017, Juniper Networks, Inc.

BEST PRACTICE: Enable these logs only to collect information when troubleshooting specific problems. Enabling these logs during normal operations can result in reduced system performance.

Related Documentation

- Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support on page 889

Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support

Problem

Description: You have collected logs on your device and need to send them to Juniper Networks Technical Support. This topic shows you how to compress the logs into a single file for each Routing Engine to more conveniently send the logs.

Solution

You can compress all the log files in the **/var/log** directories of the master and backup (if present) Routing Engines into a single **tgz** file for each Routing Engine, which enables you to send the logs to JTAC in a convenient package. You can use either the CLI or the command shell to perform these tasks; because of its ease of use, only the CLI version is shown here.

1. Access the device through the management IP address or console, typically on the master Routing Engine, RE0.

   ```
   user@host>
   ```

2. Archive and compress all the log files on RE0 and put them in **/var/tmp**.

   ```
   user@host> file archive compress source /var/log/* destination /var/tmp/re0.tgz
   /usr/bin/tar: Removing leading `/' from member names
   ```

3. Confirm that the compressed archive file has been created.

   ```
   user@host> file list /var/tmp
   baseline-config.conf
   gres-tp
   idp_license_info
   install
   jinstall-12.2-20120328.0-domestic-signed.tgz
   krt_gencfg_filter.txt
   preinstall_boot_loader.conf
   re0.tgz
   rtsdb
   sec-download
   vi.recover
   ```

   On devices with a single Routing Engine, skip to Step 10.

Copyright © 2017, Juniper Networks, Inc.

4. Log in to the backup Routing Engine, RE1, and access the CLI.

NOTE: *1* is appended to the hostname in the prompt to signify that you are on RE1.

```
user@host> request routing-engine login backup
% cli
user@host11>
```

5. Archive and compress all the log files on RE1 and put them in **/var/tmp**.

```
user@host1> file archive compress source /var/log/* destination /var/tmp/re1.tgz
/usr/bin/tar: Removing leading `/' from member names
```

6. Confirm that the compressed archive file has been created.

```
user@host1> file list /var/tmp
baseline-config.conf
gres-tp
idp_license_info
install
jinstall-12.2-20120328.0-domestic-signed.tgz
krt_gencfg_filter.txt
preinstall_boot_loader.conf
re1.tgz
rtsdb
sec-download
vi.recover
%
```

7. Exit the remote login to the backup Routing Engine to return to the master Routing Engine. Note that the previously appended *1* is removed from the hostname in the prompt to signify that you are back on RE0.

```
user@host1> exit
rlogin: connection closed

user@host1>
```

8. Copy the compressed archive file from RE1 to RE0.

```
user@host> file copy re1:/var/tmp/re1.tgz /var/tmp
```

9. Confirm the presence of the copied file.

```
user@host> file list /var/tmp
baseline-config.conf
gres-tp
idp_license_info
install
jinstall-12.2-20120328.0-domestic-signed.tgz
krt_gencfg_filter.txt
preinstall_boot_loader.conf
re0.tgz
```


Copyright © 2017, Juniper Networks, Inc.

```
rel.tgz
rtsdb
sec-download
vi.recover
%
```

10. Copy the files directly from the master Routing Engine to any local host using FTP, SCP, JWEB, or (on some devices) a mounted USB.

Related Documentation

- Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support on page 887

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 11

Configuration Statements and Operational Commands

- Configuration Statements on page 895
- Operational Commands on page 1471

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 101

Configuration Statements

- aaa-logical-system (Domain Map) on page 909
- aaa-routing-instance (Domain Map) on page 910
- abated-utilization (Address-Assignment Pools) on page 911
- abated-utilization-v6 (Address-Assignment Pools) on page 912
- accept-max-tcp-connections (System Process) on page 913
- access-identifier on page 914
- access-loop-id-local on page 914
- access-profile on page 915
- access-profile (Extensible Subscriber Services Manager) on page 916
- access-profile (Domain Map) on page 916
- access-profile (Static Subscribers) on page 917
- access-profile-name (Duplicate Accounting) on page 918
- accounting (Access Profile) on page 919
- accounting (Service Accounting) on page 920
- accounting-backup-options (Access Profile) on page 920
- accounting-order (Service Accounting) on page 921
- accounting-port on page 922
- accounting-retry (RADIUS) on page 923
- accounting-server on page 924
- accounting-session-id-format on page 924
- accounting-stop-on-access-deny on page 925
- accounting-stop-on-failure on page 926
- accounting-timeout (RADIUS) on page 927
- active-drain (Address Assignment Pools) on page 928
- active-server-group on page 929
- actual-transit-statistics (Dynamic Profiles) on page 930
- address (Diameter Peer) on page 930
- address (Diameter Transport) on page 931

Copyright © 2017, Juniper Networks, Inc.

- address-assignment (Address-Assignment Pools) on page 932
- address-change-immediate-update on page 933
- address-pool (Domain Map) on page 934
- address-protection on page 935
- adjacency-timer on page 936
- adsl-bytes on page 937
- adsl2-bytes on page 938
- adsl2-plus-bytes on page 939
- advisory-options (Traffic Shaping) on page 940
- aggregate-clients (DHCP Relay Agent) on page 941
- aggregate-clients (Static Subscribers) on page 943
- allow-bulk-leasequery (DHCP Local Server) on page 944
- allow-leasequery (DHCP Local Server) on page 945
- always-write-giaddr on page 946
- always-write-option-82 on page 947
- ancp on page 948
- ancp-speed-change-immediate-update (ANCP) on page 949
- asymmetric-lease-time (DHCP Overrides) on page 950
- asymmetric-prefix-lease-time (DHCP Overrides) on page 951
- attempts (DHCP Local Server) on page 952
- attempts (DHCP Relay Agent) on page 953
- attributes on page 954
- attributes (JSRC Attributes) on page 955
- authentication (DHCP Local Server) on page 956
- authentication (DHCP Relay Agent) on page 957
- authentication (Static Subscribers) on page 958
- authentication-order on page 959
- authentication-server on page 960
- authorization-order on page 961
- autonomous (Dynamic Router Advertisement) on page 962
- boot-file on page 962
- boot-server on page 963
- bulk-leasequery (DHCP Relay Agent) on page 964
- called-station-id (OCS Partition) on page 965
- calling-station-id-delimiter (Subscriber Management) on page 966
- calling-station-id-format (Subscriber Management) on page 967
- chap-challenge-in-request-authenticator (Subscriber Management) on page 968

Copyright © 2017, Juniper Networks, Inc.

- charging-id (OCS Partition) on page 969
- charging-service-list on page 970
- circuit-id (DHCP Relay Agent) on page 971
- circuit-id (Address-Assignment Pools) on page 973
- circuit-type (DHCP Local Server) on page 974
- circuit-type (DHCP Relay Agent) on page 975
- clear-on-abort (DHCP Local Server) on page 976
- client-accounting-algorithm on page 977
- client-authentication-algorithm on page 978
- client-discover-match (DHCP Local Server) on page 980
- client-discover-match (DHCP Relay Agent) on page 981
- client-id (DHCP Local Server) on page 982
- client-id (DHCP Relay Agent) on page 983
- client-idle-timeout on page 984
- client-idle-timeout-ingress-only on page 985
- client-negotiation-match (DHCPv6 Local Server) on page 986
- client-negotiation-match (DHCPv6 Relay Agent) on page 987
- client-session-timeout on page 988
- commit-interval on page 989
- coa-dynamic-variable-validation on page 990
- coa-immediate-update on page 990
- coa-no-override service-class-attribute on page 991
- configuration-database (Enhanced Subscriber Management) on page 991
- concurrent-data-sessions on page 992
- connect-actively on page 992
- current-hop-limit (Dynamic Router Advertisement) on page 993
- database-replication (Subscriber Session Database) on page 993
- default-action (DHCP Relay Agent Option) on page 994
- default-lifetime (Dynamic Router Advertisement) on page 995
- delay-advertise (DHCPv6) on page 996
- delay-authentication (DHCP Relay Agent) on page 998
- delay-offer (DHCPv4) on page 999
- delegated-pool (DHCP Local Server) on page 1000
- delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) on page 1001
- delimiter (DHCP Local Server) on page 1002
- delimiter (DHCP Relay Agent) on page 1004
- delimiter (Domain Map) on page 1005

Copyright © 2017, Juniper Networks, Inc.

- destination (Diameter Network Element) on page 1006
- destination-host on page 1006
- destination-host (Gx-Plus) on page 1007
- destination-host (OCS Partition) on page 1007
- destination-host (PCRF Partition) on page 1008
- destination-realm (JSRC) on page 1008
- destination-realm (Gx-Plus) on page 1009
- destination-realm (OCS Partition) on page 1010
- destination-realm (PCRF Partition) on page 1011
- dhcp-attributes (Address-Assignment Pools) on page 1012
- dhcp-local-server on page 1014
- dhcp-relay on page 1020
- dhcp-service on page 1029
- dhcpv6 (DHCP Local Server) on page 1030
- dhcpv6 (DHCP Relay Agent) on page 1034
- diameter on page 1040
- diameter-instance (JSRC) on page 1041
- diameter-instance (Diameter Applications) on page 1042
- dictionary on page 1043
- disable on page 1043
- disable (Extensible Subscriber Services Manager) on page 1044
- disable-relay on page 1045
- dne-origin (Diameter Network Element) on page 1046
- dns-server on page 1047
- dns-server-address (Dynamic Profiles) on page 1048
- domain (Domain Map) on page 1049
- domain-name (DHCP Local Server) on page 1050
- domain-name (DHCP Relay Agent) on page 1052
- domain-name (Address-Assignment Pools) on page 1053
- domain-name (Static Subscribers) on page 1054
- domain-name-server (Routing Instances and Access Profiles) on page 1055
- domain-name-server-inet (Routing Instances and Access Profiles) on page 1056
- domain-name-server-inet6 (Routing Instances and Access Profiles) on page 1057
- downstream-rate (Traffic Shaping) on page 1058
- draining (Diameter Applications) on page 1059
- draining-response-timeout (Diameter Applications) on page 1060
- drop (DHCP Relay Agent Option) on page 1061

Copyright © 2017, Juniper Networks, Inc.

- dual-stack (DHCP Relay Agent Overrides) on page 1062
- dual-stack-group (DHCP Relay Agent) on page 1063
- duplication (Access Profile) on page 1064
- duplication-filter (Access Profile) on page 1065
- duplication-vrf (Duplicate Accounting) on page 1066
- dynamic-profile (DHCP Local Server) on page 1067
- dynamic-profile (DHCP Relay Agent) on page 1068
- dynamic-profile (Domain Map) on page 1069
- dynamic-profile (Static Subscribers) on page 1070
- enable on page 1071
- enable (Enhanced Subscriber Management) on page 1071
- equals (DHCP Relay Agent) on page 1072
- ethernet-port-type-virtual on page 1073
- exceed-action on page 1073
- exclude (JSRC Attributes) on page 1074
- exclude (RADIUS) on page 1075
- external-authority on page 1080
- failover (System Process) on page 1081
- family (Address-Assignment Pools) on page 1082
- fast-lookup-filter on page 1083
- final-response-timeout (OCS Partition) on page 1084
- force-continue (OCS Partition) on page 1085
- forward-only (DHCP Relay Agent Option) on page 1086
- forward-only (DHCP Relay Agent) on page 1087
- forward-only-replies (DHCP Relay Agent) on page 1088
- forwarding (Diameter Network Element) on page 1089
- function (Diameter Network Element) on page 1090
- function (Diameter Route) on page 1091
- ggsn-address (OCS Partition) on page 1092
- ggsn-mcc-mnc (OCS Partition) on page 1093
- global (Gx-Plus) on page 1094
- global (OCS) on page 1095
- global (PCRF) on page 1096
- grace-period on page 1097
- group (DHCP Local Server) on page 1098
- group (DHCP Relay Agent) on page 1101
- group (Static Subscribers) on page 1105

Copyright © 2017, Juniper Networks, Inc.

- gsmp-syn-timeout (ANCP) on page 1106
- gsmp-syn-wait (ANCP) on page 1107
- gx-plus (Gx-Plus) on page 1108
- hardware-address on page 1109
- high-utilization (Address-Assignment Pools) on page 1109
- high-utilization-v6 (Address-Assignment Pools) on page 1110
- hold-down (Address-Assignment Pools) on page 1111
- host (Address-Assignment Pools) on page 1112
- host-name (dhcp-relay) on page 1113
- host-name (dhcpv6) on page 1114
- ietf-mode on page 1115
- ignore on page 1116
- immediate-update on page 1117
- include-ipv6 (Gx-Plus) on page 1117
- include-irb-and-l2 on page 1118
- include-l2-interface-for-irb (DHCP Relay Agent) on page 1120
- include-option-82 (DHCP Local Server) on page 1121
- interface (DHCP Local Server) on page 1122
- interface (DHCP Relay Agent) on page 1124
- interface (Dynamic Router Advertisement) on page 1126
- interface (Static Subscriber Group) on page 1127
- interface (Static Subscriber Username) on page 1128
- interface-client-limit (DHCP Local Server) on page 1129
- interface-client-limit (DHCP Relay Agent) on page 1131
- interface-delete (Subscriber Management or DHCP Client Management) on page 1132
- interface-description (DHCP Local Server) on page 1133
- interface-description (DHCP Relay Agent) on page 1134
- interface-description-format on page 1135
- interface-name (DHCP Local Server) on page 1136
- interface-name (DHCP Relay Agent) on page 1137
- interface-mib (Enhanced Subscriber Management) on page 1138
- interface-set (ANCP) on page 1139
- interface-traceoptions (DHCP) on page 1140
- interfaces (ANCP) on page 1142
- ip-address on page 1143
- ip-address-change-notify on page 1144
- ip-address-first on page 1145

Copyright © 2017, Juniper Networks, Inc.

- ip-can-type (PCRF Partition) on page 1146
- ipcp-suggest-dns-option on page 1147
- jsrc (JSRC) on page 1148
- jsrc (Access Profile) on page 1149
- jsrc-partition on page 1150
- juniper-dsl-attributes on page 1151
- layer2-unicast-replies on page 1152
- keep-incoming-circuit-id (DHCP Relay Agent) on page 1153
- keep-incoming-interface-id (DHCP Relay Agent) on page 1154
- keep-incoming-remote-id (DHCP Relay Agent) on page 1155
- leasequery (DHCP Relay Agent) on page 1156
- lease-time-threshold (DHCP Local Server and DHCP Relay Agent) on page 1157
- lease-time-validation (DHCP Local Server and DHCP Relay Agent) on page 1158
- limit on page 1158
- link (Address-Assignment Pools) on page 1159
- local (Flat-File Access Profile) on page 1160
- local-decision (PCRF Partition) on page 1161
- local-server-group (DHCP Relay Agent Option) on page 1162
- location (dhcp-relay) on page 1163
- location (dhcpv6) on page 1164
- logical-interface-unit-range on page 1165
- logical-system (Diameter Peer) on page 1166
- logical-system (Diameter Transport) on page 1167
- logical-system (Subscriber Secure Policy) on page 1168
- logical-system-name (Static Subscribers) on page 1169
- logical-system-name (DHCP Local Server) on page 1170
- logical-system-name (DHCP Relay Agent) on page 1171
- logout-response-timeout (PCRF Partition) on page 1172
- ltv-syslog-interval (System Process) on page 1173
- mac-address (DHCP Local Server) on page 1174
- mac-address (DHCP Relay Agent) on page 1175
- maintain-subscriber (Subscriber Management) on page 1176
- managed-configuration (Dynamic Router Advertisement) on page 1177
- map (Domain Map) on page 1178
- mask (Domain Map) on page 1179
- max-advertisement-interval (Dynamic Router Advertisement) on page 1180
- max-connections (DHCP Local Server) on page 1181

Copyright © 2017, Juniper Networks, Inc.

- max-data-sessions-per-subscriber on page 1182
- max-db-size (Enhanced Subscriber Management) on page 1183
- max-empty-replies (DHCP Local Server) on page 1184
- max-outstanding-requests on page 1185
- max-outstanding-requests (Diameter Applications) on page 1186
- max-pending-accounting-stops (Access Profile) on page 1187
- max-withhold-time (Access Profile) on page 1187
- maximum-discovery-table-entries on page 1188
- maximum-helper-restart-time on page 1189
- maximum-lease-time on page 1190
- maximum-subscribers on page 1191
- metric (Diameter Route) on page 1191
- metric (Domain Map) on page 1192
- min-advertisement-interval (Dynamic Router Advertisement) on page 1192
- multi-address-embedded-option-response (DHCP Local Server) on page 1193
- name-server on page 1194
- nas-identifier on page 1194
- nas-port-extended-format on page 1195
- nas-port-extended-format (Interfaces) on page 1197
- nas-port-id-delimiter (Subscriber Management) on page 1198
- nas-port-id-format (Subscriber Management) on page 1199
- nas-port-options (RADIUS Options) on page 1201
- nas-port-type (Subscriber Management) on page 1203
- nas-port-type (RADIUS Options) on page 1205
- nasreq (Diameter Application) on page 1207
- neighbor (Define ANCP) on page 1208
- neighbor-discovery-router-advertisement (Address-Assignment Pools) on page 1209
- netbios-node-type on page 1209
- network on page 1210
- network-element (Diameter Base Protocol) on page 1211
- network-services on page 1212
- no-bind-on-request (DHCP Relay Agent) on page 1214
- no-unsolicited-ra (Enhanced Subscriber Management) on page 1215
- no-vlan-interface-name on page 1216
- ocs (Diameter Applications) on page 1218
- on-demand-ip-address on page 1220
- on-link (Dynamic Router Advertisement) on page 1221

Copyright © 2017, Juniper Networks, Inc.

- option on page 1222
- option-60 (DHCP Local Server) on page 1223
- option-60 (DHCP Relay Agent) on page 1224
- option-82 (DHCP Relay Agent) on page 1225
- option-82 (DHCP Local Server Authentication) on page 1226
- option-82 (DHCP Local Server Pool Matching) on page 1227
- option-82 (Address-Assignment Pools) on page 1228
- option-match on page 1229
- option-number (DHCP Relay Agent Option) on page 1230
- options (Access Profile) on page 1231
- order on page 1233
- origin (Diameter Base Protocol) on page 1234
- other-bytes on page 1235
- other-overhead-adjust on page 1236
- other-stateful-configuration (Dynamic Router Advertisement) on page 1237
- overhead-accounting (ANCP) on page 1238
- overrides (DHCP Local Server) on page 1239
- overrides (DHCP Relay Agent) on page 1241
- overrides (Enhanced Subscriber Management) on page 1243
- override-password (Domain Map) on page 1244
- padn (Domain Map) on page 1244
- parse-direction (Domain Map) on page 1245
- parse-order (Domain Map) on page 1246
- partition on page 1247
- partition (Gx-Plus) on page 1248
- partition (NASREQ Diameter Application) on page 1249
- partition (OCS) on page 1250
- partition (PCRF) on page 1252
- password (Static Subscribers) on page 1254
- password (DHCP Local Server) on page 1255
- password (DHCP Relay Agent) on page 1257
- pcrf (Diameter Applications) on page 1259
- peer (Diameter Base Protocol) on page 1260
- peer (Diameter Network Element) on page 1261
- peer-ip-address-optional on page 1262
- peer-origin (Diameter Peer) on page 1263
- pool (Address-Assignment Pools) on page 1264

Copyright © 2017, Juniper Networks, Inc.

- pool (DHCP Local Server Overrides) on page 1265
- pool-match-order on page 1266
- port on page 1267
- port (Diameter Peer) on page 1267
- pre-ietf-mode on page 1268
- preauthentication-order (Access Profile) on page 1269
- preauthentication-port on page 1270
- preauthentication-secret on page 1271
- preauthentication-server (Access Profile) on page 1272
- preferred-lifetime (Address-Assignment Pools) on page 1273
- preferred-lifetime (Dynamic Router Advertisement) on page 1274
- prefix (DHCP Relay Agent) on page 1275
- prefix (Address-Assignment Pools) on page 1276
- prefix (Dynamic Router Advertisement) on page 1277
- priority (Diameter Peer) on page 1277
- profile (Access) on page 1278
- process-inform on page 1283
- protocols (Dynamic Profiles) on page 1285
- provisioning-order Diameter Applications) on page 1287
- proxy-mode on page 1288
- qos-adjust on page 1289
- qos-adjust-adsl on page 1290
- qos-adjust-adsl2 on page 1291
- qos-adjust-adsl2-plus on page 1292
- qos-adjust-other on page 1293
- qos-adjust-sdsl on page 1294
- qos-adjust-vdsl on page 1295
- qos-adjust-vdsl2 on page 1296
- radius (Access Profile) on page 1297
- radius-disconnect (DHCP Local Server) on page 1299
- radius-flow-tap on page 1301
- radius-options (Edit Access) on page 1302
- radius-options (Interfaces) on page 1303
- radius-server on page 1304
- range (Address-Assignment Pools) on page 1305
- rapid-commit (DHCPv6 Local Server) on page 1306
- reachable-time (Dynamic Router Advertisement) on page 1307

Copyright © 2017, Juniper Networks, Inc.

- realm-delimiter (Domain Map) on page 1307
- realm-parse-direction (Domain Map) on page 1308
- reconfigure (DHCP Local Server) on page 1309
- relay-agent-interface-id (DHCP Local Server) on page 1310
- relay-agent-interface-id (DHCPv6 Relay Agent) on page 1311
- relay-agent-interface-id (DHCPv6 Relay Agent Username) on page 1312
- relay-agent-remote-id (DHCP Local Server) on page 1313
- relay-agent-remote-id (DHCPv6 Relay Agent) on page 1314
- relay-agent-remote-id (DHCPv6 Relay Agent Username) on page 1315
- relay-agent-subscriber-id (DHCP Local Server) on page 1316
- relay-agent-subscriber-id (DHCPv6 Relay Agent) on page 1317
- relay-option (DHCP Relay Agent) on page 1318
- relay-option-vendor-specific (dhcpv6) on page 1319
- relay-option-82 on page 1320
- relay-server-group (DHCP Relay Agent Option) on page 1321
- relay-source on page 1322
- remote-id on page 1323
- remote-id (DHCP Relay Agent) on page 1324
- remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) on page 1326
- replace-ip-source-with on page 1327
- report-interface-descriptions (Edit Access) on page 1328
- report-local-rule (PCRF Partition) on page 1329
- report-resource-allocation (PCRF Partition) on page 1330
- report-successful-resource-allocation (PCRF Partition) on page 1332
- request-max-tcp-connections (System Process) on page 1333
- request network-access aaa replay pending-accounting-stops
- request network-access aaa subscriber add session-id
- request network-access aaa subscriber delete session-id
- request network-access aaa subscriber modify session-id
- request-rate on page 1340
- requested-ip-network-match (DHCP Local Server) on page 1341
- restricted-requestor (DHCP Local Server) on page 1342
- retransmit-timer (Dynamic Router Advertisement) on page 1343
- retry on page 1344
- revert-interval on page 1345
- route (Diameter Network Element) on page 1346
- router (Address-Assignment Pools) on page 1346

Copyright © 2017, Juniper Networks, Inc.

- router-advertisement (Dynamic Profiles) on page 1347
- routing-instance on page 1347
- routing-instance (Subscriber Secure Policy) on page 1348
- routing-instance (Diameter Peer) on page 1348
- routing-instance (Diameter Transport) on page 1349
- routing-instance-name (DHCP Local Server) on page 1350
- routing-instance-name (DHCP Relay Agent) on page 1352
- routing-instance-name (Static Subscribers) on page 1353
- routing-services (Enhanced Subscriber Management) on page 1354
- sdsl-bytes on page 1355
- sdsl-overhead-adjust on page 1356
- secret on page 1357
- send-acct-status-on-config-change (Access Profile) on page 1358
- send-dyn-subscription-indicator (PCRF Partition) on page 1359
- send-network-family-indicator (PCRF Partition) on page 1361
- send-origin-state-id (Diameter Applications) on page 1362
- send-release-on-delete (DHCP Relay Agent) on page 1363
- server-duid-type (DHCP Local Server) on page 1364
- server-group on page 1365
- server-identifier (Address-Assignment Pools) on page 1366
- server-id-override on page 1367
- server-response-time (DHCP Relay Agent) on page 1368
- service (Service Accounting) on page 1369
- service-activation (Access Profile) on page 1370
- service-context-id (OCS) on page 1372
- service-profile (DHCP Local Server) on page 1373
- service-profile (DHCP Relay Agent) on page 1374
- session-options on page 1375
- shmlog (Shared Memory Log) on page 1376
- sip-server-address on page 1378
- sip-server-domain-name on page 1378
- smg-service (Enhanced Subscriber Management) on page 1379
- source-address on page 1380
- stacked-vlan-ranges (RADIUS Options) on page 1381
- starts-with (DHCP Relay Agent Option) on page 1382
- static-subscribers on page 1383
- statistics (Access Profile) on page 1384

Copyright © 2017, Juniper Networks, Inc.

- statistics (Service Accounting) on page 1385
- strict (DHCP Local Server) on page 1386
- strip-domain (Domain Map) on page 1387
- strip-username (Domain Map) on page 1387
- subscriber-packet-idle-timeout on page 1388
- subscriber-management (Subscriber Management) on page 1389
- subscriber-profile on page 1390
- subscription-id-data-include (PCRF Partition) on page 1391
- subscription-id-type (PCRF Partition) on page 1393
- t1-percentage (Address-Assignment Pools) on page 1394
- t1-renewal-time (Address-Assignment Pools) on page 1395
- t2-percentage (Address-Assignment Pools) on page 1396
- t2-rebinding-time (Address-Assignment Pools) on page 1397
- target-logical-system (Domain Map) on page 1398
- target-routing-instance (Domain Map) on page 1399
- terminate-code on page 1400
- tftp-server on page 1401
- timeout (RADIUS) on page 1402
- timeout (DHCP Local Server) on page 1403
- timeout (DHCP Local Server Leasequery) on page 1404
- timeout (DHCP Relay Agent Leasequery) on page 1405
- timeout-grace (RADIUS Options) on page 1406
- token (DHCP Local Server) on page 1407
- trace (DHCP Local Server) on page 1408
- trace (DHCP Relay Agent) on page 1409
- traceoptions (ANCP) on page 1410
- traceoptions (DHCP) on page 1412
- traceoptions (Diameter Base Protocol) on page 1415
- traceoptions (Extensible Subscriber Services Manager) on page 1417
- traceoptions (Enhanced Subscriber Management) on page 1418
- traceoptions (General Authentication Service) on page 1421
- traceoptions (Static Subscribers) on page 1423
- traceoptions (Subscriber Management) on page 1425
- traceoptions (Subscriber Session Database Replication) on page 1427
- transport (Diameter Base Protocol) on page 1428
- transport (Diameter Peer) on page 1429
- traps on page 1430

Copyright © 2017, Juniper Networks, Inc.

- trigger (DHCP Local Server) on page 1431
- trigger (DHCP Relay Agent Leasequery) on page 1432
- trio-flow-offload on page 1433
- trust-option-82 on page 1434
- tunnel-profile (Domain Map) on page 1435
- underlying-interface (ANCP) on page 1435
- unique-nas-port on page 1436
- update-interval on page 1437
- update-interval (Service Accounting) on page 1438
- update-response-timeout (PCRF Partition) on page 1439
- upstream-rate (Traffic Shaping) on page 1440
- use-interface-description on page 1441
- user-name-include (OCS Partition) on page 1443
- use-option-82 on page 1445
- use-primary (DHCP Relay Agent) on page 1446
- use-vlan-id on page 1448
- user-prefix (DHCP Local Server) on page 1450
- user-prefix (DHCP Relay Agent) on page 1452
- user-prefix (Static Subscribers) on page 1453
- username-include (DHCP Local Server) on page 1454
- username-include (DHCP Relay Agent) on page 1456
- username-include (Static Subscribers) on page 1458
- valid-lifetime (Dynamic Router Advertisement) on page 1459
- valid-lifetime (Address-Assignment Pools) on page 1460
- vdsl-bytes on page 1461
- vdsl-overhead-adjust on page 1462
- vdsl2-bytes on page 1463
- vdsl2-overhead-adjust on page 1464
- vendor-specific (dhcp-relay) on page 1465
- violation-action (DHCP Local Server and DHCP Relay Agent) on page 1466
- vlan-nas-port-stacked-format on page 1467
- vlan-ranges (RADIUS Options) on page 1468
- vrf-name (Duplicate Accounting) on page 1469
- wait-for-acct-on-ack (Access Profile) on page 1469
- wins-server (Access) on page 1470

Copyright © 2017, Juniper Networks, Inc.

aaa-logical-system (Domain Map)

Syntax

```
aaa-logical-system logical-system-name {
    aaa-routing-instance routing-instance-name;
}
```

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying an AAA Logical System/Routing Instance in a Domain Map on page 169

Copyright © 2017, Juniper Networks, Inc.

aaa-routing-instance (Domain Map)

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*]

Release Information Statement introduced in Junos OS Release 10.4.
default option added in Junos OS Release 13.3.

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 (master) routing instance.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Specifying an AAA Logical System/Routing Instance in a Domain Map on page 169

Copyright © 2017, Juniper Networks, Inc.

abated-utilization (Address-Assignment Pools)

Syntax abated-utilization *percentage*;

Hierarchy Level [edit access address-assignment],
[edit routing-instances *routing-instance-name* address-assignment]

Release Information Statement introduced in Junos OS Release 11.2.

Description 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.

Options ***percentage***—Threshold below which an SNMP trap clear is generated.

Range: 1 through 98

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pool Usage Threshold Traps on page 588

Copyright © 2017, Juniper Networks, Inc.

abated-utilization-v6 (Address-Assignment Pools)

Syntax abated-utilization-v6 *percentage*;

Hierarchy Level [edit access address-assignment],
[edit routing-instances *routing-instance-name* address-assignment]

Release Information Statement introduced in Junos OS Release 11.2.

Description 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.

Options ***percentage***—Threshold below which an SNMP trap clear is generated.

Range: 1 through 98

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pool Usage Threshold Traps on page 588

Copyright © 2017, Juniper Networks, Inc.

accept-max-tcp-connections (System Process)

Syntax accept-max-tcp-connections *max-tcp-connections*;

Hierarchy Level [edit system processes dhcp-service]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the number of simultaneous TCP connections that the DHCP server can use for bulk leasequery operations.

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.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

access-identifier

Syntax access-identifier *identifier-string*;

Hierarchy Level [edit protocols ancp interfaces *interface-name*],
[edit protocols ancp interfaces interface-set]

Release Information Statement introduced in Junos OS Release 9.4.

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.

Related Documentation
- Configuring the ANCP Agent on page 648
- Associating an Access Node with Subscribers for ANCP Agent Operations on page 651

access-loop-id-local

Syntax access-loop-id-local;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

access-profile

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*]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

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.

Related Documentation

- Attaching Access Profiles on page 124
- Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 332
- *Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*

Copyright © 2017, Juniper Networks, Inc.

access-profile (Extensible Subscriber Services Manager)

Syntax access-profile *access-profile-name;*

Hierarchy Level **[edit system services extensible-subscriber-services]**

Release Information Statement introduced in Junos OS Release 15.1.

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.

access-profile (Domain Map)

Syntax access-profile *profile-name*;

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation

- Specifying an Access Profile in a Domain Map on page 167

Copyright © 2017, Juniper Networks, Inc.

access-profile (Static Subscribers)

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*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Global Access Profile on page 805
- Specifying the Static Subscriber Group Access Profile on page 810

Copyright © 2017, Juniper Networks, Inc.

access-profile-name (Duplicate Accounting)

Syntax access-profile-name [*profile-name*];

Hierarchy Level [edit access profile *profile-name* accounting duplication-vrf]

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.

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.

Related Documentation

- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

accounting (Access Profile)

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*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring Per-Subscriber Session Accounting on page 105
- Understanding RADIUS Accounting Duplicate Reporting on page 99

Copyright © 2017, Juniper Networks, Inc.

accounting (Service Accounting)

Syntax

```
accounting {
  statistics (time | volume-time);
  update-interval minutes;
}
```

Hierarchy Level [edit access profile *profile-name* service]

Release Information Statement introduced in Junos OS Release 14.2R1 for MX Series routers.

Description Define the service accounting configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Service Interim Accounting on page 112
- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155

accounting-backup-options (Access Profile)

Syntax

```
accounting-backup-options {
  max-pending-accounting-stops number;
  max-withhold-time hold-time;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 13.1.

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. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Configuring Back-up Options for RADIUS Accounting on page 111

Copyright © 2017, Juniper Networks, Inc.

accounting-order (Service Accounting)

Syntax accounting-order (activation-protocol | local | radius);

Hierarchy Level [edit access profile *profile-name* service]

Release Information Statement introduced in Junos OS Release 11.4.
local option added in Junos OS Release 17.1.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Service Accounting with JSRC on page 794
- Service Accounting with JSRC on page 793
- *Configuring Service Accounting in Local Flat Files*

accounting-port

Syntax accounting-port *port-number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS).
Statement introduced on Junos OS without ELS in the following releases:

- Junos OS Release 12.3 for EX Series switches: Release 12.3R10.
- Junos OS Release 14.1X53 for EX Series switches: Release 14.1X53-D25.
- Junos OS Release 15.1 for EX Series switches: Release 15.1R4.

Description 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.

Options ***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

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring RADIUS System Accounting*
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2017, Juniper Networks, Inc.

accounting-retry (RADIUS)

Syntax accounting-retry *number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the number of times the router 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.

Options *number*—Number of retry attempts.

Range: 0 through 100

Default: 0 (disabled)

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2017, Juniper Networks, Inc.

accounting-server

Syntax accounting-server [*ip-address*];

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

Options ***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.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

accounting-session-id-format

Syntax accounting-session-id-format (decimal | description);

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the format the router or switch uses to identify the accounting session.

Default **decimal**

Options **decimal**—Use the decimal format.

description—Use the generic format, in the form: **jnpr** *interface-specifier*:*subscriber-session-id*.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

accounting-stop-on-access-deny

Syntax accounting-stop-on-access-deny;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

accounting-stop-on-failure

Syntax accounting-stop-on-failure;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

 Copyright © 2017, Juniper Networks, Inc.

accounting-timeout (RADIUS)

Syntax accounting-timeout *seconds*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure how long the local router 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.

Options *seconds*—Duration of timeout period.

Default: 0 (disabled)

Range: 0 through 1000 seconds

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2017, Juniper Networks, Inc.

active-drain (Address Assignment Pools)

Syntax active-drain;

Hierarchy Level [edit access address-assignment pool *pool-name*]

Release Information Statement introduced in Junos OS Release 17.2 on MX Series.

Description 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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring DHCP Local Address Pool Rapid Drain on page 591
- Configuring Address-Assignment Pools on page 585
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

active-server-group

Syntax active-server-group *server-group-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 group *group-name* 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* forwarding-options group *group-name*],
[edit logical-systems *logical-system-name* forwarding-options group *group-name* 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 logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* dhcpv6],
[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 group *group-name*],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group *group-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Apply a DHCP relay agent configuration to the named group of DHCP server addresses. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

A group-specific configuration overrides a global option.

Options ***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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups on page 335
- Configuring Group-Specific DHCP Relay Options on page 290
- **dhcp-relay on page 1020**

Copyright © 2017, Juniper Networks, Inc.

actual-transit-statistics (Dynamic Profiles)

Syntax actual-transit-statistics

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces unit "$junos-interface-unit"]

Release Information Support at the **[edit dynamic-profiles *profile-name* interfaces unit "$junos-interface-unit"]** hierarchy level introduced in Junos OS Release 15.1R3.

Description Enable **actual-transit-statistics** to engage subscriber accounting. Statistics provide packet-level and byte-level counts for all transmit and receive traffic forwarded across the specified interface. The counts do not include filtered traffic, drops, discards, or control traffic.

Default Disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)*
- *Dynamic Profiles Overview*

address (Diameter Peer)

Syntax address *ip-address*;

Hierarchy Level [edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the IP address for a Diameter remote peer.

Options ***ip-address***—IP address of remote Diameter peer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

address (Diameter Transport)

Syntax address *ip-address*;

Hierarchy Level [edit diameter transport *transport-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure the source (local) IP address for the Diameter local transport connection.

Options ***ip-address***—IP address of remote Diameter peer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723

Copyright © 2017, Juniper Networks, Inc.

address-assignment (Address-Assignment Pools)

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;
      }
      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]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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 *pool-name*—Name assigned to an address-assignment pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- *Configuring an Address-Assignment Pool for L2TP LNS with Inline Services*

address-change-immediate-update

Syntax address-change-immediate-update;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the router to send an Interim-Accounting message to the RADIUS server immediately after on-demand IPv4 allocation and de-allocation.

Changes to this setting take effect for new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Default This functionality is disabled by default.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes on page 471
- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461

Copyright © 2017, Juniper Networks, Inc.

address-pool (Domain Map)

Syntax address-pool *pool-name*;

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation

- Specifying an Address Pool in a Domain Map on page 168

Copyright © 2017, Juniper Networks, Inc.

address-protection

Syntax address-protection;

Hierarchy Level [edit access],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* access]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Duplicate IPv6 Prefix Protection for Router Advertisement on page 487

Copyright © 2017, Juniper Networks, Inc.

- Configuring Duplicate IPv4 Address Protection for AAA on page 595

adjacency-timer

Syntax adjacency-timer *seconds*;

Hierarchy Level [edit protocols ancp],
[edit protocols ancp neighbor *ip-address*]

Release Information Statement introduced in Junos OS Release 9.4.

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.

Related Documentation

- Configuring the ANCP Agent on page 648
- Specifying the Interval Between ANCP Adjacency Messages on page 652
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

adsl-bytes

Syntax adsl-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 17.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

adsl2-bytes

Syntax adsl2-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 17.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

adsl2-plus-bytes

Syntax adsl2-plus-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 17.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

advisory-options (Traffic Shaping)

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*]

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.

Description

Specify a recommended shaping rate to be applied to downstream or upstream traffic on an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 683
- Configuring the ANCP Agent on page 648
- *Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS*

Copyright © 2017, Juniper Networks, Inc.

aggregate-clients (DHCP Relay Agent)

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*]

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.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- dhcp-relay on page 1020
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- Configuring Group-Specific DHCP Relay Options on page 290

Copyright © 2017, Juniper Networks, Inc.

aggregate-clients (Static Subscribers)

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*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

This statement is not supported for IP demux subscriber interfaces.

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.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers on page 806

Copyright © 2017, Juniper Networks, Inc.

allow-bulk-leasequery (DHCP Local Server)

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 ...]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Enable DHCP local server to listen for, and respond to bulk leasequery requests received on TCP connections.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

allow-leasequery (DHCP Local Server)

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 ...]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Enable DHCP local server to listen on the UDP socket and to respond to leasequery requests received on UDP connections.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

always-write-giaddr

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]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

Related Documentation
- Extended DHCP Relay Agent Overview on page 240
- dhcp-relay on page 1020

Copyright © 2017, Juniper Networks, Inc.

always-write-option-82

Syntax always-write-option-82;

Hierarchy Level
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Extended DHCP Relay Agent Overview on page 240

Copyright © 2017, Juniper Networks, Inc.

ancp

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;
    }
```


Copyright © 2017, Juniper Networks, Inc.

}

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure Junos OS ANCP agent features.

The remaining statements are explained separately. See CLI Explorer.

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.

Related Documentation

- Configuring the ANCP Agent on page 648

ancp-speed-change-immediate-update (ANCP)

Syntax ancp-speed-change-immediate-update;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 13.3.

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.

Related Documentation

- Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications on page 692
- Configuring Per-Subscriber Session Accounting on page 105
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

asymmetric-lease-time (DHCP Overrides)

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...]

Release Information Statement introduced in Junos OS Release 17.1.

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.

Related Documentation

- DHCP Asymmetric Leasing Overview on page 261
- Configuring DHCP Asymmetric Leasing on page 263

Copyright © 2017, Juniper Networks, Inc.

asymmetric-prefix-lease-time (DHCP Overrides)

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...]

Release Information Statement introduced in Junos OS Release 17.1.

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.

Related Documentation
- DHCP Asymmetric Leasing Overview on page 261
- Configuring DHCP Asymmetric Leasing on page 263

Copyright © 2017, Juniper Networks, Inc.

attempts (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

Copyright © 2017, Juniper Networks, Inc.

- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 310

attempts (DHCP Relay Agent)

Syntax attempts *number-of-attempts*;

Hierarchy Level [edit forwarding-options dhcp-relay (leasequery | bulk-leasequery)],
[edit forwarding-options dhcp-relay dhcpv6 (leasequery | bulk-leasequery)],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* ...],
[edit routing-instances *routing-instance-name* ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the number of times the DHCP relay agent attempts to send DHCP leasequery or 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.

Options *number-of-attempts*—Number of times DHCP relay agent attempts to send the leasequery or bulk leasequery message.

For the DHCP bulk leasequery messages, DHCP leasequery messages, DHCPv6 leasequery messages:

Range: 1 through 10

Default: 6

For the DHCPv6 bulk leasequery messages:

Range: 1 through 720

Default: 360

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

attributes

Syntax

```
attributes {
  exclude {
    ...
  }
  ignore {
    framed-ip-netmask;
    input-filter;
    logical-system-routing-instance;
    output-filter;
  }
}
```

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Specify how the router or switch processes RADIUS attributes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43

Copyright © 2017, Juniper Networks, Inc.

attributes (JSRC Attributes)

Syntax

```
attributes {
  exclude {
    user-name [ authorization-request | provisioning-request ];
    }
  }
}
```

Hierarchy Level [edit access profile *profile-name* jsrc]

Release Information Statement introduced in Junos OS Releases 14.2.

Description Specify how the router processes Diameter attributes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Excluding AVPs from Diameter Messages for JSRC on page 791
- Understanding JSRC-SAE Interactions on page 785

Copyright © 2017, Juniper Networks, Inc.

authentication (DHCP Local Server)

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;
  }
}
```

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 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 ...]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

authentication (DHCP Relay Agent)

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;
  }
}
```

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 ...]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1020
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

authentication (Static Subscribers)

Syntax

```
authentication {
  password password-string;
  username-include {
    domain-name domain-name;
    interface;
    logical-system-name;
    routing-instance-name;
    user-prefix user-prefix-string;
  }
}
```

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*]

Release Information

Statement introduced in Junos OS Release 9.6.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Authentication Password on page 807
- Configuring the Static Subscriber Group Authentication Password on page 811

Copyright © 2017, Juniper Networks, Inc.

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
none option added in Junos OS Release 11.2.
nasreq option added in Junos OS Release 16.1.

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.

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 the client using NASREQ authentication services.
- **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.
- **password**—Verify the client using the information configured at the **[edit access profile *profile-name* client *client-name*]** hierarchy level.
- **radius**—Verify the client using RADIUS authentication services.

NOTE: Subscriber access management does not support the password option, and authentication fails when no method (none) is specified.

Default: password

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring CHAP Authentication with RADIUS*
- Specifying the Authentication and Accounting Methods for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

- *Configuring Access Profiles for L2TP or PPP Parameters*

authentication-server

Syntax authentication-server [*ip-address*];

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.

Description 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.

Options ***ip-address***—IPv4 address.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116

 Copyright © 2017, Juniper Networks, Inc.

authorization-order

Syntax authorization-order [*authorization-methods*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
nasreq option added in Junos OS Release 16.1.

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 DHCP subscriber can access the router or switch. When you include this option, AAA ignores any configured authentication order settings. This option is ignored for 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.

Related Documentation

- Configuring JSRC on page 787
- Authorizing Subscribers with JSRC on page 789

Copyright © 2017, Juniper Networks, Inc.

autonomous (Dynamic Router Advertisement)

Syntax (autonomous | no-autonomous);

Hierarchy Level [edit dynamic-profiles *profile-name* protocols router-advertisement interface *interface-name* prefix *prefix*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

boot-file

Syntax boot-file *filename*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description 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.

Options *filename*—Location of the boot file on the boot server. The filename can include a pathname.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- boot-server on page 963
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

boot-server

Syntax boot-server (*address* | *hostname*);

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description 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.

Options ***address***—IPv4 address of a boot server.

hostname—Fully qualified hostname of a boot server.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- boot-file on page 962
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

bulk-leasequery (DHCP Relay Agent)

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* ...]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Configure parameters that the DHCP relay agent uses when sending DHCP bulk leasequery messages to obtain lease information from the DHCP servers in the logical system/routing instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

called-station-id (OCS Partition)

Syntax called-station-id *station-name*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

calling-station-id-delimiter (Subscriber Management)

Syntax calling-station-id-delimiter *delimiter-character*;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

Description 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.

Default The hash (**#**) character.

Options ***delimiter-character***—Character to use for the delimiter. You must enclose the delimiter character in quotation marks (“ ”).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Calling-Station-ID with Additional Attributes on page 66

Copyright © 2017, Juniper Networks, Inc.

calling-station-id-format (Subscriber Management)

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]

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.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Calling-Station-ID with Additional Attributes on page 66

chap-challenge-in-request-authenticator (Subscriber Management)

Syntax
chap-challenge-in-request-authenticator;

Hierarchy Level
[edit access profile *profile-name* radius options]

Release Information
Statement introduced in Junos OS Release 15.1.

Description
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.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

charging-id (OCS Partition)

Syntax charging-id *number*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

charging-service-list

Syntax charging-service-list ocs;

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

circuit-id (DHCP Relay Agent)

Syntax

```
circuit-id {
  include-irb-and-l2;
  include-l2-interface-for-irb;
  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]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
include-irb-and-l2 , **no-vlan-interface-name**, and **use-vlan-id** options added in Junos OS Release 14.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring Option 82 Information on page 350

circuit-id (Address-Assignment Pools)

Syntax circuit-id *value* range *named-range*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes option-match option-82],
[edit access protocol-attributes *attribute-set-name* option-match option-82]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure the address-assignment pool named-range to use for a particular option 82 Agent Circuit ID value.

Options ***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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

circuit-type (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

circuit-type (DHCP Relay Agent)

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 ...]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

clear-on-abort (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 311

Copyright © 2017, Juniper Networks, Inc.

client-accounting-algorithm

Syntax client-accounting-algorithm (direct | round-robin);

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS for EX Series switches Release 13.2X50-D10.

Description Configure the access method the router uses to access RADIUS accounting servers.

Default direct

Options **direct**—Use the direct method.

round-robin—Use the round-robin method.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

client-authentication-algorithm

Syntax client-authentication-algorithm (direct | round-robin);

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure the method that the authenticator uses to access RADIUS authentication servers when there are multiple servers configured. Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The router or switch, acting as the authenticator, waits for a response from the server before sending another request.

When there are multiple RADIUS server connections configured for a client, the authenticator attempts to reach the different servers in the order that they are configured. If there is no response from the first RADIUS server, the authenticator attempts to reach the next RADIUS server. This process repeats until the client is either granted access or there are no more configured servers.

If the **direct** method is configured, the authenticator always treats the first server in the list as the primary server. The authenticator moves on to the second server only if the attempt to reach the first server fails. If the **round-robin** method is configured, the server chosen first will be rotated based on which server was used last. The first server in the list is treated as a primary for the first authentication request, but for the second request, the second server configured is treated as primary, and so on. With this method, all of the configured servers receive roughly the same number of requests on average so that no single server has to handle all of the requests.

NOTE: The round-robin access method is not recommended for use with EX Series switches.

Default The **direct** option is the default.

Options

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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

client-discover-match (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
incoming-interface option added in Junos OS Release 13.3.

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.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Overriding Default DHCP Local Server Configuration Settings on page 321
- DHCP Auto Logout Overview on page 313
- Allowing Only One DHCP Client Per Interface on page 298

Copyright © 2017, Juniper Networks, Inc.

client-discover-match (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
incoming-interface option added in Junos OS Release 13.3.

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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324
- DHCP Auto Logout Overview on page 313
- Allowing Only One DHCP Client Per Interface on page 298

Copyright © 2017, Juniper Networks, Inc.

client-id (DHCP Local Server)

Syntax client-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]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify that the DHCPv6 Client-ID option (option 1) in the client PDU name 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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

client-id (DHCP Relay Agent)

Syntax client-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 ...],
[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 ...]

Release Information
Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify that the client ID 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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

client-idle-timeout

Syntax client-idle-timeout *minutes*;

Hierarchy Level [edit access profile *profile-name* session-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description 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.

Options ***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

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138
- Removing Inactive Dynamic Subscriber VLANs on page 139

client-idle-timeout-ingress-only

Syntax
client-idle-timeout-ingress-only;

Hierarchy Level
[edit access profile *profile-name* session-options]

Release Information
Statement introduced in Junos OS Release 16.2.

Description
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.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138
- Removing Inactive Dynamic Subscriber VLANs on page 139

Copyright © 2017, Juniper Networks, Inc.

client-negotiation-match (DHCPv6 Local Server)

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Automatically Logging Out DHCPv6 Clients on page 318
- Extended DHCP Local Server Overview on page 234
- Allowing Only One DHCP Client Per Interface on page 298

Copyright © 2017, Juniper Networks, Inc.

client-negotiation-match (DHCPv6 Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Automatically Logging Out DHCPv6 Clients on page 318
- Extended DHCP Relay Agent Overview on page 240
- Allowing Only One DHCP Client Per Interface on page 298

Copyright © 2017, Juniper Networks, Inc.

client-session-timeout

Syntax client-session-timeout *minutes*;

Hierarchy Level [edit access profile *profile-name* session-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description 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.

Options ***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

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138

commit-interval

Syntax
commit-interval *interval*;

Hierarchy Level
[edit system services extensible-subscriber-services]

Release Information
Statement introduced in Junos OS Release 15.1.

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.

Related Documentation
- *Op Script Overview*

Copyright © 2017, Juniper Networks, Inc.

coa-dynamic-variable-validation

Syntax coa-dynamic-variable-validation;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description 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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- RADIUS Server Options for Subscriber Access on page 4
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

coa-immediate-update

Syntax coa-immediate-update;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 10.0.

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.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring Per-Subscriber Session Accounting on page 105

Copyright © 2017, Juniper Networks, Inc.

coa-no-override service-class-attribute

Syntax coa-no-override service-class-attribute;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring Per-Subscriber Session Accounting on page 105

configuration-database (Enhanced Subscriber Management)

Syntax
```
configuration-database {
  max-db-size size;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers with MPCs

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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Junos OS Enhanced Subscriber Management on page 575

Copyright © 2017, Juniper Networks, Inc.

concurrent-data-sessions

Syntax concurrent-data-sessions *max-session-number*;

Hierarchy Level [edit services service-set *services-set-name* subscriber-profile *profile-name* enable *service-name*]

Release Information Statement introduced in Junos OS Release 11.4.

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.

connect-actively

Syntax

```
connect-actively {
  port port-number;
  transport transport-name;
}
```

Hierarchy Level [edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Define the destination port and transport connection used to establish active connections to Diameter peer.

The remaining statements are explained separately. See CLI Explorer.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

current-hop-limit (Dynamic Router Advertisement)

Syntax current-hop-limit *number*;

Hierarchy Level [edit dynamic-profiles protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

database-replication (Subscriber Session Database)

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]

Release Information Statement introduced in Junos OS Release 9.3.

Description Define operations for subscriber management session database replication processes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

Copyright © 2017, Juniper Networks, Inc.

default-action (DHCP Relay Agent Option)

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 ...]

Release Information

Statement introduced in Junos OS Release 12.3.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

default-lifetime (Dynamic Router Advertisement)

Syntax default-lifetime *seconds*;

Hierarchy Level [edit protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *max-advertisement-interval (Protocols IPv6 Neighbor Discovery)*
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

delay-advertise (DHCPv6)

Syntax

```
delay-advertise {
  based-on {
    option 15 {
      equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        not-equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        starts-with {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
      }
    option 16 {
      equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        not-equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        starts-with {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        }
        option 18 {
        equals {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
          not-equals {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
          starts-with {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
          }
        option 37 {
          equals {
            ascii <ascii-name>;
              hexadecimal <hex-name>;
            }
            not-equals {
            ascii <ascii-name>;
              hexadecimal <hex-name>;
```


Copyright © 2017, Juniper Networks, Inc.

```
            }
            starts-with {
            ascii <ascii-name>;
              hexadecimal <hex-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]

Release Information Statement introduced in Junos OS Release 16.1.

Description Delay the DHCP-OFFER sent to the customers. Delay can be set between 1 to 30 seconds.

Default If no delay time is set, default delay time of 3 seconds will be used.

Options **delay-time <*time in seconds*>**—Time delay between solicit and advertise. Time delay can be set between 1 to 30 seconds.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

delay-authentication (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 13.3.

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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

delay-offer (DHCPv4)

Syntax

```
delay-offer {
  based-on {
    option 60 {
      equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        not-equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        starts-with {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
      }
      option 77 {
      equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        not-equals {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        starts-with {
        ascii <ascii-name>;
          hexadecimal <hex-name>;
        }
        }
        option 82 {
        equals {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
          not-equals {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
          starts-with {
          ascii <ascii-name>;
            hexadecimal <hex-name>;
          }
        }
      }
  }
```

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]

Copyright © 2017, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 16.1.

Description Delay the DHCP-OFFER sent to the customers. Delay can be set between 1 to 30 seconds.

Default If no delay time is set, default delay time of 3 seconds will be used.

Options **delay-time <*time in seconds*>**—Time delay between discover and offer. Time delay can be set between 1 to 30 seconds.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

delegated-pool (DHCP Local Server)

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 ...]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation

- Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation on page 490
- Overriding Default DHCP Local Server Configuration Settings on page 321

Copyright © 2017, Juniper Networks, Inc.

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

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]

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure the router to override the default behavior when the router receives a DHCPv4 Discover or DHCPv6 Solicit message while in a bound state. By default, DHCP maintains the existing client entry when DHCP receives a new 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.

You use the **delete-binding-on-renegotiation** statement to override the default action, and to specify that DHCP tear down the existing matching client entry and to process the message as a new client entry.

NOTE: In Junos OS releases prior to 15.1, the DHCPv6 local server and relay agent use the opposite default behavior, and tear down the existing client entry when receiving a Solicit message while in a bound state. Starting in Junos OS release 15.1, the DHCPv6 local server and DHCPv6 relay agent default behavior is the same as DHCPv4 local server and relay agent—in all models, the default behavior is to maintain the existing client entry.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- DHCP Behavior When Renegotiating While in Bound State on page 329

Copyright © 2017, Juniper Networks, Inc.

delimiter (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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)

Copyright © 2017, Juniper Networks, Inc.

NOTE: When you include the interface-description 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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

delimiter (DHCP Relay Agent)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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)

Copyright © 2017, Juniper Networks, Inc.

NOTE: When you include the interface-description 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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

delimiter (Domain Map)

Syntax
delimiter [*delimiter-character*];

Hierarchy Level
[edit access domain]

Release Information
Statement introduced in Junos OS Release 10.4.

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.

Related Documentation
- Specifying Domain and Realm Name Delimiters on page 172
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

destination (Diameter Network Element)

Syntax destination realm *realm-name* <host *hostname*>;

Hierarchy Level [edit diameter network-element *element-name* forwarding route *dne-route-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

destination-host

Syntax destination-host *hostname*

Hierarchy Level [edit jsrc partition *partition-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788

Copyright © 2017, Juniper Networks, Inc.

destination-host (Gx-Plus)

Syntax destination-host *hostname*;

Hierarchy Level [edit access gx-plus partition *partition-name*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Related Documentation
- Configuring Gx-Plus on page 736
- Configuring the Gx-Plus Partition on page 737

destination-host (OCS Partition)

Syntax destination-host *ocs-hostname*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation
- Configuring the OCS Partition on page 769
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

destination-host (PCRF Partition)

Syntax destination-host *pcrf-hostname*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

destination-realm (JSRC)

Syntax destination-realm *realm*

Hierarchy Level [edit jsrc partition *partition-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788

Copyright © 2017, Juniper Networks, Inc.

destination-realm (Gx-Plus)

Syntax destination-realm *realm*;

Hierarchy Level [edit access gx-plus partition *partition-name*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Related Documentation

- Configuring Gx-Plus on page 736
- Configuring the Gx-Plus Partition on page 737

Copyright © 2017, Juniper Networks, Inc.

destination-realm (OCS Partition)

Syntax destination-realm *ocs-realm-name*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

destination-realm (PCRF Partition)

Syntax destination-realm *pcrf-realm-name*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

dhcp-attributes (Address-Assignment Pools)

Syntax

```
dhcp-attributes {
  boot-file filename;
  boot-server (address | hostname);
  dns-server [ ipv6-address ];
  domain-name domain-name;
  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*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

- Configuring Address-Assignment Pools on page 585
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254

Copyright © 2017, Juniper Networks, Inc.

dhcp-local-server

Syntax

```
dhcp-local-server {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    access-profile profile-name;
    authentication {
      ...
    }
    duplicate-clients incoming-interface;
    group group-name {
      access-profile profile-name;
      authentication {
        ...
      }
      interface interface-name {
        access-profile profile-name;
        exclude;
        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;
            }
          }
```


Copyright © 2017, Juniper Networks, Inc.

```
        }
        overrides {
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        delegated-pool;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    route-suppression;
    server-duid-type type;
    service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
```

Copyright © 2017, Juniper Networks, Inc.

```
        delegated-pool;
        include-option-82 {
            forcerenew;
            nak;
        }
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    reconfigure {
        attempts attempt-count;
        clear-on-abort;
        strict;
        timeout timeout-value;
        token token-value;
        trigger {
            radius-disconnect;
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
forward-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;
        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);
```


Copyright © 2017, Juniper Networks, Inc.

```
                holddown-interval milliseconds;
              }
            }
          }
          overrides {
            asymmetric-lease-time seconds;
            client-discover-match (option60-and-option82 | incoming-interface);
            include-option-82 {
              forcerenew;
              nak;
            }
            interface-client-limit number;
            process-inform {
              pool pool-name;
            }
            protocol-attributes attribute-set-name;
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        overrides {
          asymmetric-lease-time seconds;
          client-discover-match (option60-and-option82 | incoming-interface);
          include-option-82 {
            forcerenew;
            nak;
          }
          interface-client-limit number;
          process-inform {
            pool pool-name;
          }
          protocol-attributes attribute-set-name;
        }
        requested-ip-network-match subnet-mask
        route-suppression;
        service-profile dynamic-profile-name;
      }
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
```

```
        }
      }
    }
    overrides {
      asymmetric-lease-time seconds;
      client-discover-match <option60-and-option82 | incoming-interface>;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      protocol-attributes attribute-set-name;
    }
    pool-match-order {
      external-authority;
      ip-address-first;
      option-82;
    }
    reconfigure {
      attempts attempt-count;
      clear-on-abort;
      strict;
      timeout timeout-value;
      token token-value;
      trigger {
        radius-disconnect;
      }
    }
    requested-ip-network-match subnet-mask;
    route-suppression;
    service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],
[edit logical-systems *logical-system-name* system services],
[edit routing-instances *routing-instance-name* system services],
[edit system services]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

dhcp-relay

Syntax

```
dhcp-relay {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  bulk-leasequery {
    attempts number-of-attempts;
    timeout seconds;
    trigger automatic;
  }
  dhcpv6 {
    access-profile profile-name;
    active-server-group server-group-name;
    }
    authentication {
      password password-string;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description (device-interface | logical-interface);
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    bulk-leasequery {
      attempts number-of-attempts;
      timeout seconds;
      trigger automatic;
    }
    duplicate-clients incoming-interface;
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
```


Copyright © 2017, Juniper Networks, Inc.

```
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
}
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        ...
    }
    dual-stack-group dual-stack-group-name {
        access-profile profile-name;
        authentication {
            ... authentication-configuration
        }
        dynamic-profile profile-name {
            ... dynamic-profile-configuration
        }
        relay-agent-interface-id {
            ... relay-agent-interface-id-configuration
        }
        relay-agent-remote-id {
            ... relay-agent-remote-id-configuration
        }
        service-profile dynamic-profile-name;
    }
    dynamic-profile 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 {
            ...
        }
        exclude;
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
```

```
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  }
  lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
      route-suppression;
      service-profile dynamic-profile-name;
    }
  }
  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 {
    ...
  }
  relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
  }
  relay-option {
```


Copyright © 2017, Juniper Networks, Inc.

```
        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 action;
      route-suppression;
      service-profile dynamic-profile-name;
    }
    leasequery {
      attempts number-of-attempts;
      timeout seconds;
    }
    lease-time-validation {
      lease-time-threshold seconds;
      violation-action action;
    }
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode(automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
        route-suppression;
        service-profile dynamic-profile-name;
      }
    }
    no-snoop;
    overrides {
      allow-snooped-clients;
      asymmetric-lease-time seconds;
```

Copyright © 2017, Juniper Networks, Inc.

```
        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 {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    relay-agent-remote-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            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 action;
    route-suppression;
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
}
dual-stack-group dual-stack-group-name {
    access-profile profile-name;
    authentication {
        ... authentication-configuration
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
    dynamic-profile profile-name {
        ... dynamic-profile-configuration
    }
    relay-agent-interface-id {
        ... relay-agent-interface-id-configuration
    }
    relay-agent-remote-id {
        ... relay-agent-remote-id-configuration
    }
    service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile 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;
```

Copyright © 2017, Juniper Networks, Inc.

```
                }
                detection-time {
                  threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
              }
            }
          }
          overrides {
            ...
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        overrides {
          ...
        }
        relay-option {
          ...
        }
        relay-option-82 {
          ...
        }
        route-suppression:
        service-profile dynamic-profile-name;
      }
      leasequery {
        attempts number-of-attempts;
        timeout seconds;
      }
      lease-time-validation {
        lease-time-threshold seconds;
        violation-action action;
      }
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
        }
```


Copyright © 2017, Juniper Networks, Inc.

```
}
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 action;
```

Copyright © 2017, Juniper Networks, Inc.

```
  route-suppression:
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  server-response-time seconds;
service-profile dynamic-profile-name;
```

Hierarchy Level

[edit forwarding-options],
[edit logical-systems *logical-system-name* forwarding-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options],
[edit routing-instances *routing-instance-name* forwarding-options]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455
- DHCP Relay Proxy Overview on page 243
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

dhcp-service

Syntax

```
dhcp-service {
  accept-max-tcp-connections max-tcp-connections;
  dhcp-snooping-file(local_pathname | remote_URL ) {
    write-interval interval;
  }
  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;
  }
  ltv-syslog-interval seconds;
  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]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1 for MX Series routers.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)*

Copyright © 2017, Juniper Networks, Inc.

dhcpv6 (DHCP Local Server)

Syntax

```
dhcpv6 {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      logical-system-name;
      mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  duplicate-clients incoming-interface;
  group group-name {
    access-profile profile-name;
    authentication {
    ...
    interface interface-name {
      access-profile profile-name;
      exclude;
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
```


Copyright © 2017, Juniper Networks, Inc.

```
            delete-binding-on-renegotiation;
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
              pool pool-name;
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        overrides {
          asymmetric-lease-time seconds;
          asymmetric-prefix-lease-time seconds;
          client-negotiation-match incoming-interface;
          delegated-pool;
          delete-binding-on-renegotiation;
          interface-client-limit number;
          multi-address-embedded-option-response;
          process-inform {
            pool pool-name;
          }
          protocol-attributes attribute-set-name;
          rapid-commit;
        }
        route-suppression;
        service-profile dynamic-profile-name;
      }
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delegated-pool;
```

```
    delete-binding-on-renegotiation;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
      pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
    reconfigure {
      attempts attempt-count;
      clear-on-abort;
      strict;
      timeout timeout-value;
      token token-value;
      trigger {
        radius-disconnect;
      }
    }
  }
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
  requested-ip-network-match subnet-mask;
  route-suppression;
  server-duid-type type;
  service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

dhcpv6 (DHCP Relay Agent)

Syntax

```
dhcpv6 {
    access-profile profile-name;
    active-server-group server-group-name;
    }
    authentication {
      password password-string;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description (device-interface | logical-interface);
        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;
      }
    }
    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;
    }
    group group-name {
      access-profile profile-name;
      active-server-group server-group-name;
      authentication {
        ...
      }
      dual-stack-group dual-stack-group-name {
        access-profile profile-name;
        authentication {
          ... authentication-configuration
        }
        dynamic-profile profile-name {
          ... dynamic-profile-configuration
        }
        relay-agent-interface-id {
```


Copyright © 2017, Juniper Networks, Inc.

```
      ... relay-agent-interface-id-configuration
    }
    relay-agent-remote-id {
      ... relay-agent-remote-id-configuration
    }
    service-profile dynamic-proflle-name;
  }
  dynamic-profile 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 {
      ...
    }
    exclude;
    overrides {
      allow-snooped-clients;
      asymmetric-lease-time seconds;
      asymmetric-prefix-lease-time seconds;
      client-negotiation-match incoming-interface;
      delay-authentication;
      delete-binding-on-renegotiation;
      dual-stack dual-stack-group-name;
      interface-client-limit number;
      no-allow-snooped-clients;
      no-bind-on-request;
      relay-source interface-name;
      send-release-on-delete;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  }
  lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
```

Copyright © 2017, Juniper Networks, Inc.

```
            detection-time {
              threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
          route-suppression;
          service-profile dynamic-profile-name;
        }
      }
      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 {
        ...
      }
      relay-agent-remote-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
      }
      relay-option {
        option-number option-number;
        default-action {
          drop;
          forward-only;
          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 action;
      route-suppression;
      service-profile dynamic-profile-name;
    }
    leasequery {
      attempts number-of-attempts;
```


Copyright © 2017, Juniper Networks, Inc.

```
    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;
        }
        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 {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
```

Copyright © 2017, Juniper Networks, Inc.

```
    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 action;
  route-suppression;
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  server-response-time seconds;
  service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

Release Information

Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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. See CLI Explorer.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1020
- DHCPv6 Relay Agent Overview on page 455
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

diameter

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;
  }
  transport transport-name {
    address;
    logical-system logical-system-name <routing-instance routing-instance-name>;
    routing-instance routing-instance-name;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the Diameter base protocol for subscriber management.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723

Copyright © 2017, Juniper Networks, Inc.

diameter-instance (JSRC)

Syntax diameter-instance *instance-name*

Hierarchy Level [edit jsrc partition *partition-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788

Copyright © 2017, Juniper Networks, Inc.

diameter-instance (Diameter Applications)

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*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
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.

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.

Related Documentation

- Configuring Gx-Plus on page 736
- Configuring the Gx-Plus Partition on page 737
- Configuring the OCS Partition on page 769
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

dictionary

Syntax dictionary *dictionary-path;*

Hierarchy Level **[edit system services extensible-subscriber-services]**

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- show extensible-subscriber-services dictionary on page 1712
- show extensible-subscriber-services dictionary attributes on page 1716
- request services extensible-subscriber-services reload-dictionary on page 1536
- show extensible-subscriber-services dictionary services on page 1719
- Understanding the Dictionary File on page 226

disable

Syntax disable *service-name*;

Hierarchy Level [edit services service-set *services-set-name* subscriber-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Copyright © 2017, Juniper Networks, Inc.

disable (Extensible Subscriber Services Manager)

Syntax disable

Hierarchy Level **[edit system processes extensible-subscriber-services]**

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- restart extensible-subscriber-services on page 1546

 Copyright © 2017, Juniper Networks, Inc.

disable-relay

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]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Extended DHCP Relay Agent Overview on page 240

Copyright © 2017, Juniper Networks, Inc.

dne-origin (Diameter Network Element)

Syntax dne-origin realm *realm-name* <host *hostname*>;

Hierarchy Level [edit diameter network-element *element-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring Diameter Network Elements on page 726
- Configuring Diameter on page 723

Copyright © 2017, Juniper Networks, Inc.

dns-server

Syntax dns-server *ipv6-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description 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.

Options ***ipv6-address***—IPv6 address of a DNS server.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

dns-server-address (Dynamic Profiles)

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*]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation

- DNS Resolver for IPv6 DNS Overview on page 604
- Configuring a DNS Server Address for IPv6 Hosts on page 604

Copyright © 2017, Juniper Networks, Inc.

domain (Domain Map)

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;
    padn destination-address {
      mask destination-mask;
      metric route-metric;
    }
    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]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure a domain map, which is used to map access options and session parameters for subscriber sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Domain Map on page 164

Copyright © 2017, Juniper Networks, Inc.

domain-name (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

domain-name (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

domain-name (Address-Assignment Pools)

Syntax
domain-name *domain-name*;

Hierarchy Level
[edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information
Statement introduced in Junos OS Release 9.0.

Description
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.

Options
domain-name—Name of the domain.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

domain-name (Static Subscribers)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

domain-name-server (Routing Instances and Access Profiles)

Syntax domain-name-server *dns-address*;

Hierarchy Level [edit access];
[edit access profile]

Release Information Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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 statement. That is, the server with the address configured with the domain-name-server-inet 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.

Related Documentation

- Configuring DNS Name Server Addresses for Subscriber Management on page 600
- DNS Name Server Address Overview on page 599

Copyright © 2017, Juniper Networks, Inc.

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

Syntax domain-name-server-inet *dns-address*;

Hierarchy Level [edit access],
[edit access profile]

Release Information Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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 statement. That is, the server with the address configured with the domain-name-server-inet takes precedence over a server configured with the domain-name-server 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.

Related Documentation

- Configuring DNS Name Server Addresses for Subscriber Management on page 600
- DNS Name Server Address Overview on page 599

Copyright © 2017, Juniper Networks, Inc.

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

Syntax domain-name-server-inet6 *dns-address*;

Hierarchy Level [edit access],
[edit access profile]

Release Information Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Related Documentation

- Configuring DNS Name Server Addresses for Subscriber Management on page 600
- DNS Name Server Address Overview on page 599

Copyright © 2017, Juniper Networks, Inc.

downstream-rate (Traffic Shaping)

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]

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.

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.

Related Documentation
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 683
- Configuring the ANCP Agent on page 648
- *Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS*

Copyright © 2017, Juniper Networks, Inc.

draining (Diameter Applications)

Syntax draining;

Hierarchy Level [edit access pcrf partition *partition-name*],
[edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

draining-response-timeout (Diameter Applications)

Syntax draining-response-timeout *seconds*;

Hierarchy Level [edit access pcrf partition *partition-name*],
[edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

drop (DHCP Relay Agent Option)

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 ...]

Release Information

Statement introduced in Junos OS Release 12.3.

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.

Related Documentation

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

dual-stack (DHCP Relay Agent Overrides)

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 ...]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation
- Single-Session DHCP Dual-Stack Overview on page 422
- Configuring Single-Session DHCP Dual-Stack Support on page 423

Copyright © 2017, Juniper Networks, Inc.

dual-stack-group (DHCP Relay Agent)

Syntax

```
dual-stack-group dual-stack-group-name {
  access-profile profile-name;
  authentication {
    ... authentication-configuration
  }
  dynamic-profile profile-name {
    ... dynamic-profile-configuration
  }
  relay-agent-interface-id {
    ... relay-agent-interface-id-configuration
  }
  relay-agent-remote-id {
    ... relay-agent-remote-id-configuration
  }
  service-profile dynamic-proflle-name;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP dual stack, and names the dual stack group.

The group is assigned to each leg of the DHCP dual-stack with the dual-stack statement in the overrides stanza. When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement in the configuration.

Related Documentation

- Single-Session DHCP Dual-Stack Overview on page 422
- Configuring Single-Session DHCP Dual-Stack Support on page 423

Copyright © 2017, Juniper Networks, Inc.

duplication (Access Profile)

Syntax duplication;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Understanding RADIUS Accounting Duplicate Reporting on page 99

Copyright © 2017, Juniper Networks, Inc.

duplication-filter (Access Profile)

Syntax duplication-filter (interim-duplicated | interim-original) <exclude-attributes>;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting on page 101
- Configuring Per-Subscriber Session Accounting on page 105

Copyright © 2017, Juniper Networks, Inc.

duplication-vrf (Duplicate Accounting)

Syntax

```
duplication-vrf {
  access-profile-name profile-name;
  vrf-name vrf-name;
}
```

Hierarchy Level [edit access profile *profile-name* accounting]

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.)

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

dynamic-profile (DHCP Local Server)

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 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 ...]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- Configuring a Default Subscriber Service on page 248

Copyright © 2017, Juniper Networks, Inc.

dynamic-profile (DHCP Relay Agent)

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 ...]

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1020
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- Grouping Interfaces with Common DHCP Configurations on page 287
- Configuring a Default Subscriber Service on page 248

Copyright © 2017, Juniper Networks, Inc.

dynamic-profile (Domain Map)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation

- Specifying a Dynamic Profile in a Domain Map on page 168

Copyright © 2017, Juniper Networks, Inc.

dynamic-profile (Static Subscribers)

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*]

Release Information

Statement introduced in Junos OS Release 9.6.

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 statement is explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Specifying the Static Subscriber Global Dynamic Profile on page 805
- Specifying the Static Subscriber Group Dynamic Profile on page 810

Copyright © 2017, Juniper Networks, Inc.

enable

Syntax

```
enable service-name {
  concurrent-data-sessions max-session-number;
}
```

Hierarchy Level [edit services service-set *services-set-name* subscriber-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable the service name for the subscriber profile.

Options ***service-name***—Name of the enabled service.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

enable (Enhanced Subscriber Management)

Syntax enable;

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description 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.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Junos OS Enhanced Subscriber Management Overview on page 573

Copyright © 2017, Juniper Networks, Inc.

equals (DHCP Relay Agent)

Syntax

```
equals (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 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 ...]

Release Information

Statement introduced in Junos OS Release 12.3.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

ethernet-port-type-virtual

Syntax ethernet-port-type-virtual;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description 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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

exceed-action

Syntax
```
exceed-action {
  drop;
  syslog;
}
```

Hierarchy Level [edit services service-set *services-set-name* subscriber-profile *profile-name* max-data-sessions-per-subscriber]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Copyright © 2017, Juniper Networks, Inc.

exclude (JSRC Attributes)

Syntax

```
exclude {
  user-name [ authorization-request | provisioning-request ];
}
```

Hierarchy Level [edit access profile *profile-name* jsrc attributes]

Release Information Statement introduced in Junos OS Release 14.2.

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.

Related Documentation

- Excluding AVPs from Diameter Messages for JSRC on page 791
- Understanding JSRC-SAE Interactions on page 785

Copyright © 2017, Juniper Networks, Inc.

exclude (RADIUS)

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-on | accounting-off | accounting-start |
      accounting-stop ];
    acc-loop-remote-id [ access-request | accounting-on | accounting-off | accounting-start
      | accounting-stop ];
    accounting-authentic [ accounting-on | accounting-off ];
    accounting-delay-time [ accounting-on | accounting-off ];
    accounting-session-id [ access-request | accounting-on | accounting-off | accounting-stop
      ];
    accounting-terminate-cause [ accounting-off ];
    acct-tunnel-connection [ 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 ];
    chap-challenge [ access-request ];
    chargeable-user-identity [ access-request ];
    class [ accounting-start | accounting-stop ];
    cos-shaping-rate [ accounting-start | accounting-stop ];
    dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    dhcp-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 ];
    event-timestamp [ accounting-on | accounting-off | 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-ip-address [ accounting-start | accounting-stop ];
    framed-ip-netmask [ accounting-start | accounting-stop ];
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets [ accounting-stop ];
    input-gigawords [ accounting-stop ];
    interface-description [ 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 ];
```

Copyright © 2017, Juniper Networks, Inc.

```
    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-on | accounting-off | 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 ];
    pppoe-description [ access-request | accounting-start | accounting-stop ];
    tunnel-assignment-id [ accounting-start | accounting-stop ];
    tunnel-client-auth-id [ accounting-start | accounting-stop ];
    tunnel-client-endpoint [ accounting-start | accounting-stop ];
    tunnel-medium-type [ accounting-start | accounting-stop ];
    tunnel-server-auth-id [ accounting-start | accounting-stop ];
    tunnel-server-endpoint [ accounting-start | accounting-stop ];
    tunnel-type [ accounting-start | accounting-stop ];
    upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
    virtual-router [ access-request | accounting-start | accounting-stop ]
}
```

Hierarchy Level [edit access profile *profile-name* radius attributes]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Options **downstream-calculated-qos-rate**, **dsl-forum-attributes**, and **upstream-calculated-qos-rate** added in Junos OS Release 11.4.
Options **cos-shaping-rate** and **filter-id** added in Junos OS Release 13.2.
Option **pppoe-description** added in Junos OS Release 14.2.
Option **virtual-router** added in Junos OS Release 15.1.
Options **first-relay-ipv4-address** and **first-relay-ipv6-address** added in Junos OS Release 16.1.
Options **acc-loop-encap** and **acc-loop-remote-id** added in Junos OS Release 16.1R4.

Description Configure the router or switch to exclude the specified attributes from the specified type of RADIUS message.

Not all attributes are available in all types of RADIUS messages. By default, the router or switch includes the specified attributes in RADIUS Access-Request, Acct-On, Acct-Off, Acct-Start, and Acct-Stop messages.

NOTE: If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

Copyright © 2017, Juniper Networks, Inc.

Options RADIUS attribute type—RADIUS attribute, Juniper Networks (vendor ID 4874) VSA number and name, or DSL Forum (vendor ID 3561) VSA number and name:

- **acc-aggr-cir-id-asc**—Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- **acc-aggr-cir-id-bin**—Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- **acc-loop-cir-id**—Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- **acc-loop-encap**—Juniper Networks VSA 26-183, Acc-Loop-Encap.
- **acc-loop-remote-id**—Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- **accounting-authentic**—RADIUS attribute 45, Acct-Authentic.
- **accounting-delay-time**—RADIUS attribute 41, Acct-Delay-Time.
- **accounting-session-id**—RADIUS attribute 44, Acct-Session-Id.
- **accounting-terminate-cause**—RADIUS attribute 49, Acct-Terminate-Cause.
- **acct-tunnel-connection**—RADIUS attribute 68, Acct-Tunnel-Connection.
- **act-data-rate-dn**—Juniper Networks VSA 26-114, Act-Data-Rate-Dn
- **act-data-rate-up**—Juniper Networks VSA 26-113, Act-Data-Rate-Up
- **act-interlv-delay-dn**—Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn
- **act-interlv-delay-up**—Juniper Networks VSA 26-124, Act-Interlv-Delay-Up
- **att-data-rate-dn**—Juniper Networks VSA 26-118, Att-Data-Rate-Dn
- **att-data-rate-up**—Juniper Networks VSA 26-117, Att-Data-Rate-Up
- **called-station-id**—RADIUS attribute 30, Called-Station-Id.
- **calling-station-id**—RADIUS attribute 31, Calling-Station-Id.
- **chargeable-user-identity**—RADIUS attribute 89, Chargeable-User-Identity.
- **class**—RADIUS attribute 25, Class.
- **cos-shaping-rate**—Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- **dhcp-gi-address**—Juniper Networks VSA 26-57, DHCP-GI-Address.
- **dhcp-mac-address**—Juniper Networks VSA 26-56, DHCP-MAC-Address.
- **dhcp-options**—Juniper Networks VSA 26-55, DHCP-Options.
- **downstream-calculated-qos-rate**—Juniper Networks VSA 26-141
- **dsl-forum-attributes**—DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*
- **dsl-line-state**—Juniper Networks VSA 26-127, DSL-Line-State
- **dsl-type**—Juniper Networks VSA 26-128, DSL-Type
- **event-timestamp**—RADIUS attribute 55, Event-Timestamp.
- **filter-id**—RADIUS attribute 11, Filter-Id.

Copyright © 2017, Juniper Networks, Inc.

- **first-relay-ipv4-address** —Juniper Networks VSA 26-187, DHCP-First-Relay-IPv4-Address.
- **first-relay-ipv6-address** —Juniper Networks VSA 26-188, DHCP-First-Relay-IPv6-Address.
- **framed-ip-address**—RADIUS attribute 8, Framed-IP-Address.
- **framed-ip-netmask**—RADIUS attribute 9, Framed-IP-Netmask.
- **input-filter**—Juniper Networks VSA 26-10, Ingress-Policy-Name.
- **input-gigapackets**—Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- **input-gigawords**—RADIUS attribute 52, Acct-Input-Gigawords.
- **interface-description**—Juniper Networks VSA 26-53, Interface-Desc.
- **l2tp-rx-connect-speed**—Juniper Networks VSA 26-163, Rx-Connect-Speed
- **l2tp-tx-connect-speed**—Juniper Networks VSA 26-162, Tx-Connect-Speed
- **max-data-rate-dn**—Juniper Networks VSA 26-120, Max-Data-Rate-Dn
- **max-data-rate-up**—Juniper Networks VSA 26-119, Max-Data-Rate-Up
- **max-interlv-delay-dn**—Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn
- **max-interlv-delay-up**—Juniper Networks VSA 26-123, Max-Interlv-Delay-Up
- **min-data-rate-dn**—Juniper Networks VSA 26-116, Min-Data-Rate-Dn
- **min-data-rate-up**—Juniper Networks VSA 26-115, Min-Data-Rate-Up
- **min-lp-data-rate-dn**—Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn
- **min-lp-data-rate-up**—Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up
- **nas-identifier**—RADIUS attribute 32, NAS-Identifier.
- **nas-port**—RADIUS attribute 5, NAS-Port.
- **nas-port-id**—RADIUS attribute 87, NAS-Port-Id.
- **nas-port-type**—RADIUS attribute 61, NAS-Port-Type.
- **output-filter**—Juniper Networks VSA 26-11, Egress-Policy-Name.
- **output-gigapackets**—Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- **output-gigawords**—RADIUS attribute 53, Acct-Output-Gigawords.
- **pppoe-description**—Juniper Networks VSA 26-24, PPPoE-Description.
- **tunnel-assignment-id**—RADIUS attribute 82, Tunnel-Assignment-ID.
- **tunnel-client-auth-id**—RADIUS attribute 90. Tunnel-Client-Auth-ID.
- **tunnel-client-endpoint**—RADIUS attribute 66, Tunnel-Client-Endpoint.
- **tunnel-medium-type**—RADIUS attribute 65, Tunnel-Medium-Type.
- **tunnel-server-auth-id**—RADIUS attribute 91, Tunnel-Server-Auth-ID.
- **tunnel-server-endpoint**—RADIUS attribute 67, Tunnel-Server-Endpoint.

Copyright © 2017, Juniper Networks, Inc.

- **tunnel-type**—RADIUS attribute 64, Tunnel-Type.
- **upstream-calculated-qos-rate**—Juniper Networks VSA 26-142
- **virtual-router**—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.

Related Documentation

- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

external-authority

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]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation

- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- Extended DHCP Local Server Overview on page 234
- Address-Assignment Pools Overview on page 583

Copyright © 2017, Juniper Networks, Inc.

failover (System Process)

Syntax failover (alternate-media | other-routing-engine);

Hierarchy Level [edit system processes *process-name*]

Release Information Statement introduced before Junos OS Release 7.4.

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 mastership; 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.

Related Documentation

- *When a Software Process Fails*
- *processes*

Copyright © 2017, Juniper Networks, Inc.

family (Address-Assignment Pools)

Syntax

```
family family {
  dhcp-attributes {
    [protocol-specific attributes]
  }
  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*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

fast-lookup-filter

Syntax fast-lookup-filter;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall family *family*],
[edit firewall family *family-name* filter *filter-name*],
[edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name*]

Release Information Statement introduced in Junos OS Release 13.3R3 for MX 240, MX 480, MX 960, MX 2010, and MX 2020 routers with MPC5E, MPC5EQ, and MPC6E MPCs. Support for the **next-header** firewall match condition is available in Junos OS Release 13.3R6 and later.

Description The **fast-lookup-filter** is available for the inet and inet6 protocol families for both static and dynamic profiles. Junos installs firewall filters created under this hierarchy to the accelerated filter block available in the MPC5E, MPC5EQ, and MPC6E MPCs, which provides enhanced performance.

Juniper recommends that you use the **payload-protocol** term rather than the **next-header** term when configuring a firewall filter with match conditions for IPv6 traffic. Although either can be used, **payload-protocol** provides the more reliable match condition because it uses the actual payload protocol to find a match, whereas **next-header** simply takes whatever appears in the first header following the IPv6 header, which may or may not be the actual protocol. In addition, if **next-header** is used with IPv6, the accelerated filter block lookup process is bypassed and the standard filter used instead.

See *Firewall Filter Match Conditions for IPv6 Traffic* for more information about firewall filters and terms.

The filter block hardware supports up to 4095 firewall filters, each of which can support up to 255 terms, to a system maximum of 8000 terms. Firewall instances from the same firewall block can also be attached to multiple interfaces.

NOTE: If both **fast-filter-optimization** and **fast-lookup-filter** are configured at the same time, **fast-lookup-filter** will take precedence (that is, any terms under **fast-filter-optimization** will be ignored).

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

Copyright © 2017, Juniper Networks, Inc.

final-response-timeout (OCS Partition)

Syntax final-response-timeout *seconds*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

force-continue (OCS Partition)

Syntax force-continue;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

forward-only (DHCP Relay Agent Option)

Syntax forward-only;

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 ...]

Release Information Statement introduced in Junos OS Release 12.3.

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.

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 configured overrides options.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

forward-only (DHCP Relay Agent)

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 ...]

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

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.

Related Documentation

- DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs on page 339
- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs on page 340

Copyright © 2017, Juniper Networks, Inc.

forward-only-replies (DHCP Relay Agent)

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 ...]

Release Information Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

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.

Related Documentation

- DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs on page 339
- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs on page 340

Copyright © 2017, Juniper Networks, Inc.

forwarding (Diameter Network Element)

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*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Define the criteria that specify which destinations are reachable through the Diameter network element.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

function (Diameter Network Element)

Syntax function [*function-name*];

Hierarchy Level [edit diameter network-element *element-name*]

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.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

function (Diameter Route)

Syntax function *function-name* <partition *partition-name*>;

Hierarchy Level [edit diameter network-element *element-name* forwarding route *dne-route-name*]

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.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

ggsn-address (OCS Partition)

Syntax ggsn-address *address*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

ggsn-mcc-mnc (OCS Partition)

Syntax ggsn-mcc-mnc *ggsn-mcc-mnc*;

Hierarchy Level [edit access ocs partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

global (Gx-Plus)

Syntax

```
global {
  include-ipv6;
  max-outstanding-requests number;
}
```

Hierarchy Level [edit access gx-plus]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure global attributes for the Gx-Plus application.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Gx-Plus on page 736

Copyright © 2017, Juniper Networks, Inc.

global (OCS)

Syntax

```
global {
  service-context-id service-context ;
}
```

Hierarchy Level [edit access ocs]

Release Information Statement introduced in Junos OS Release 16.2.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring OCS Global Parameters on page 776
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

global (PCRF)

Syntax

```
global {
  rule-param avp-code;
}
```

Hierarchy Level [edit access pcrf]

Release Information Statement introduced in Junos OS Release 16.2.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

grace-period

Syntax grace-period *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description 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.

Options ***seconds***—Number of seconds the lease is retained.

Range: 0 through 4,294,967,295 seconds

Default: 0 (no grace period)

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

group (DHCP Local Server)

Syntax

```
group group-name {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      relay-agent-interface-id
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
  interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
      asymmetric-lease-time seconds;
      asymmetric-prefix-lease-time seconds;
      client-discover-match <option60-and-option82>;
      client-negotiation-match incoming-interface;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
```


Copyright © 2017, Juniper Networks, Inc.

```
            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-discover-match <option60-and-option82>;
      client-negotiation-match incoming-interface;
      delegated-pool;
      delete-binding-on-renegotiation;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      protocol-attributes attribute-set-name;
      rapid-commit;
    }
    reconfigure {
      attempts attempt-count;
      clear-on-abort;
      strict;
      timeout timeout-value;
      token token-value;
      trigger {
        radius-disconnect;
      }
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Copyright © 2017, Juniper Networks, Inc.

Options

group-name—Name of the group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Grouping Interfaces with Common DHCP Configurations on page 287
- Using External AAA Authentication Services with DHCP on page 281
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

Copyright © 2017, Juniper Networks, Inc.

group (DHCP Relay Agent)

Syntax

```
group group-name {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      logical-system-name;
      mac-address;
      option-60;
      option-82 [circuit-id] [remote-id];
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  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;
        }
```

Copyright © 2017, Juniper Networks, Inc.

```
        }
    }
    overrides {
        ...
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82>;
    client-negotiation-match incoming-interface;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
```


Copyright © 2017, Juniper Networks, Inc.

```
        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
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information

Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

Options

group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- **dhcp-relay on page 1020**
- Extended DHCP Relay Agent Overview on page 240
- *Understanding the Extended DHCP Relay Agent for EX Series Switches*
- *Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)*
- Configuring Group-Specific DHCP Relay Options on page 290
- Grouping Interfaces with Common DHCP Configurations on page 287
- Using External AAA Authentication Services with DHCP on page 281
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

Copyright © 2017, Juniper Networks, Inc.

group (Static Subscribers)

Syntax

```
group group-name {
  access-profile profile-name;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
  }
  authentication {
    password password-string;
    username-include {
      domain-name domain-name;
      interface;
      logical-system-name;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  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]

Release Information

Statement introduced in Junos OS Release 9.6.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Creating a Static Subscriber Group on page 809

gsmp-syn-timeout (ANCP)

Syntax gsmp-syn-timeout *seconds*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 13.3.

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.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring the ANCP Agent to Learn ANCP Partition IDs on page 654

Copyright © 2017, Juniper Networks, Inc.

gsmp-syn-wait (ANCP)

Syntax gsmp-syn-wait;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 13.3

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.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring the ANCP Agent to Learn ANCP Partition IDs on page 654

Copyright © 2017, Juniper Networks, Inc.

gx-plus (Gx-Plus)

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]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure the Gx-Plus application to interact with a PCRF to authorize and provision subscribers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Gx-Plus on page 736

Copyright © 2017, Juniper Networks, Inc.

hardware-address

Syntax hardware-address *mac-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) host *hostname*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the MAC address of the client. This is the hardware address that identifies the client on the network.

Options ***mac-address***—MAC address of the client.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Address-Assignment Pools on page 585

high-utilization (Address-Assignment Pools)

Syntax high-utilization *percentage*;

Hierarchy Level [edit access address-assignment],
[edit routing-instances *routing-instance-name* address-assignment]

Release Information Statement introduced in Junos OS Release 11.2.

Description 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.

Options ***percentage***—Percentage used to generate a trap.
Range: 2 through 99

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pool Usage Threshold Traps on page 588

Copyright © 2017, Juniper Networks, Inc.

high-utilization-v6 (Address-Assignment Pools)

Syntax high-utilization-v6 *percentage*;

Hierarchy Level [edit access address-assignment],
[edit routing-instances *routing-instance-name* address-assignment]

Release Information Statement introduced in Junos OS Release 11.2.

Description 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.

Options ***percentage***—Percentage used to generate a trap.

Range: 2 through 99

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pool Usage Threshold Traps on page 588

Copyright © 2017, Juniper Networks, Inc.

hold-down (Address-Assignment Pools)

Syntax hold-down;

Hierarchy Level [edit access address-assignment pool *pool-name*]

Release Information Statement introduced in Junos OS Release 16.1.

Description 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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- Configuring Address-Assignment Pool Hold-Down on page 590

Copyright © 2017, Juniper Networks, Inc.

host (Address-Assignment Pools)

Syntax

```
host hostname {
  hardware-address mac-address;
  ip-address ip-address;
}
```

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6)]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure a static binding for the specified client.

Options ***hostname***—Name of the client.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

host-name (dhcp-relay)

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]

Release Information Command introduced in Junos OS Release 16.2.

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. The 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.

Copyright © 2017, Juniper Networks, Inc.

host-name (dhcpv6)

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]

Release Information Command introduced in Junos OS Release 16.2.

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. The Enterprise ID is 2636. The hostname is option-data 1 (the location is option-data 2). The DHCPv6 relay strips the sub-option 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.

Copyright © 2017, Juniper Networks, Inc.

ietf-mode

Syntax ietf-mode

Hierarchy Level [edit protocols ancp neighbor *ip-address*]

Release Information Statement introduced in Junos OS Release 9.5.

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.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

ignore

Syntax

```
ignore {
  dynamic-iflset-name;
  framed-ip-netmask;
  input-filter;
  logical-system-routing-instance;
  output-filter;
}
```

Hierarchy Level [edit access profile *profile-name* radius attributes]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. By default, the router or switch processes the attributes it receives from the external server.

Options

dynamic-iflset-name—Ignore Interface-Set/Dynamic-Iflset-Name (VSA 26-130).

framed-ip-netmask—Ignore Framed-IP-Netmask (RADIUS attribute 9).

input-filter—Ignore Ingress-Policy-Name (VSA 26-10).

logical-system-routing-instance—Ignore Virtual-Router (VSA 26-1).

output-filter—Ignore Egress-Policy-Name (VSA 26-11).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 43
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

immediate-update

Syntax immediate-update;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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.

Related Documentation
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring Per-Subscriber Session Accounting on page 105

include-ipv6 (Gx-Plus)

Syntax include-ipv6;

Hierarchy Level [edit access gx-plus global]

Release Information Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Related Documentation
- Configuring Gx-Plus Global Attributes on page 738
- Configuring Gx-Plus on page 736

Copyright © 2017, Juniper Networks, Inc.

include-irb-and-l2

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]

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.

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.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Including a Textual Description in DHCP Options on page 354
- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring DHCPv6 Relay Agent Options on page 456

Copyright © 2017, Juniper Networks, Inc.

include-l2-interface-for-irb (DHCP Relay Agent)

Syntax include-l2-interface-for-irb;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* relay-agent-interface-id],
[edit forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit forwarding-options dhcp-relay group *group-name* relay-option-82 circuit-id],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* ...],
[edit routing-instances *routing-instance-name* ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description 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.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270
- Using DHCP Relay Agent Option 82 Information on page 349
- Including a Prefix in DHCP Options on page 352
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458

Copyright © 2017, Juniper Networks, Inc.

include-option-82 (DHCP Local Server)

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]

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

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.

Related Documentation

- DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs on page 339
- Configuring DHCP Message Exchange Between DHCP Server and Clients in Different VRFs on page 340

Copyright © 2017, Juniper Networks, Inc.

interface (DHCP Local Server)

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;
    interface-client-limit number;
    rapid-commit;
  }
  service-profile dynamic-profile-name;
  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 ...]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Options **upto** and **exclude** introduced in Junos OS Release 9.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Grouping Interfaces with Common DHCP Configurations on page 287
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

interface (DHCP Relay Agent)

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;
  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 ...]

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.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Grouping Interfaces with Common DHCP Configurations on page 287
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

interface (Dynamic Router Advertisement)

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]

Release Information Statement introduced in Junos OS Release 10.1.

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. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

interface (Static Subscriber Group)

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*]

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.

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.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Creating a Static Subscriber Group on page 809

Copyright © 2017, Juniper Networks, Inc.

interface (Static Subscriber Username)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

interface-client-limit (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Default No limit

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Specifying the Maximum Number of DHCP Clients Per Interface on page 297
- Overriding Default DHCP Local Server Configuration Settings on page 321

Copyright © 2017, Juniper Networks, Inc.

interface-client-limit (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description 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.

Default No limit

Options *number*—Maximum number of clients allowed.
Range: 1 through 500,000

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- dhcp-relay on page 1020
- Extended DHCP Relay Agent Overview on page 240
- Configuring Group-Specific DHCP Relay Options on page 290
- Overriding the Default DHCP Relay Configuration Settings on page 324

interface-delete (Subscriber Management or DHCP Client Management)

Syntax
interface-delete;

Hierarchy Level
[edit system services subscriber-management maintain-subscriber]

Release Information
Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302

Copyright © 2017, Juniper Networks, Inc.

interface-description (DHCP Local Server)

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 ...]

Release Information Statement introduced in Junos OS Release 17.1.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

interface-description (DHCP Relay Agent)

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 ...]

Release Information Statement introduced in Junos OS Release 17.1.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

interface-description-format

Syntax

```
interface-description-format {
  exclude-adapter;
  exclude-sub-interface;
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Options **exclude-adapter** and **exclude-sub-interface** introduced in Junos OS Release 10.4.

Description Specify the information that is excluded from the interface description that the device passes to RADIUS for inclusion in the RADIUS attribute 87 (NAS-Port-Id). By default, the device includes both the subinterface and the adapter in the interface description.

Options

exclude-adapter—Exclude the adapter from the interface description.

exclude-sub-interface—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.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- RADIUS Server Options for Subscriber Access on page 4

Copyright © 2017, Juniper Networks, Inc.

interface-name (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

 Copyright © 2017, Juniper Networks, Inc.

interface-name (DHCP Relay Agent)

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 ...]

Release Information
Statement introduced in Junos OS Release 11.4
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

interface-mib (Enhanced Subscriber Management)

Syntax interface-mib;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Enable representation of the Interfaces MIB for the specified dynamic interface.

BEST PRACTICE: 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.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Junos OS Enhanced Subscriber Management Overview on page 573

Copyright © 2017, Juniper Networks, Inc.

interface-set (ANCP)

Syntax

```
interface-set interface-set-name {
    access-identifier identifier-string;
    underlying-interface underlying-interface-name;
}
```

Hierarchy Level [edit protocols ancp interfaces]

Release Information Statement introduced in Junos OS Release 9.4.

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. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the ANCP Agent on page 648
- Associating an Access Node with Subscribers for ANCP Agent Operations on page 651

Copyright © 2017, Juniper Networks, Inc.

interface-traceoptions (DHCP)

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]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*****k** to specify KB, ***size*****m** to specify MB, or ***size*****g** to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Tracing Extended DHCP Operations for Specific Interfaces on page 834

Copyright © 2017, Juniper Networks, Inc.

interfaces (ANCP)

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]

Release Information Statement introduced in Junos OS Release 9.4.

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. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the ANCP Agent on page 648
- Associating an Access Node with Subscribers for ANCP Agent Operations on page 651

Copyright © 2017, Juniper Networks, Inc.

ip-address

Syntax ip-address *ip-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet host *hostname*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the reserved IP address assigned to the client.

Options ***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.

Related Documentation

- Configuring Address-Assignment Pools on page 585
- Configuring Static Address Assignment on page 594

Copyright © 2017, Juniper Networks, Inc.

ip-address-change-notify

Syntax

```
ip-address-change-notify {
  message message;
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

Description For on-demand address allocation for dual-stack PPP subscribers, specify that the BNG includes the Unisphere-Ipv4-release-control VSA 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.

Options ***message***—VSA message.
Range: Up to 32 characters.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes on page 471

Copyright © 2017, Juniper Networks, Inc.

ip-address-first

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]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

Related Documentation

- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- Extended DHCP Local Server Overview on page 234
- Address-Assignment Pools Overview on page 583

Copyright © 2017, Juniper Networks, Inc.

ip-can-type (PCRF Partition)

Syntax ip-can-type *number*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

ipcp-suggest-dns-option

Syntax ipcp-suggest-dns-option;

Hierarchy Level [edit access group-profile *group-profile-name* ppp ppp-options],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit *logical-unit-number* ppp-options],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number* ppp-options]

Release Information Statement introduced in Junos OS Release 16.1.

Description Configure the router to prompt Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for terminated PPPoE and LNS subscribers. You can configure this for dynamic or static PPPoE subscribers, dynamic or static LNS subscribers, and in an LNS group profile.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses on page 602
- *Configuring the Group Profile for Defining L2TP Attributes*
- *Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile*
- *Configuring Dynamic Authentication for PPP Subscribers*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

Copyright © 2017, Juniper Networks, Inc.

jsrc (JSRC)

Syntax

```
jsrc {
  partition partition-name {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm-name;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure JSRC to interact with an SAE in an SRC environment to authorize and provision subscribers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring JSRC on page 787

Copyright © 2017, Juniper Networks, Inc.

jsrc (Access Profile)

Syntax

```
jsrc {
  attributes {
    exclude {
      user-name [ authorization-request | provisioning-request ];
      }
    }
  }
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify JSRC settings in an access profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring JSRC on page 787
- Authorizing Subscribers with JSRC on page 789
- Understanding JSRC-SAE Interactions on page 785
- Excluding AVPs from Diameter Messages for JSRC on page 791

Copyright © 2017, Juniper Networks, Inc.

jsrc-partition

Syntax jsrc-partition *partition-name*;

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788

Copyright © 2017, Juniper Networks, Inc.

juniper-dsl-attributes

Syntax juniper-dsl-attributes;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure AAA to add Juniper Networks DSL 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.

Default The Juniper Networks DSL 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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 691
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

layer2-unicast-replies

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]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- dhcp-relay on page 1020

Copyright © 2017, Juniper Networks, Inc.

keep-incoming-circuit-id (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Using DHCP Relay Agent Option 82 Information on page 349

Copyright © 2017, Juniper Networks, Inc.

keep-incoming-interface-id (DHCP Relay Agent)

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],

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458

 Copyright © 2017, Juniper Networks, Inc.

keep-incoming-remote-id (DHCP Relay Agent)

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],

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation
- Using DHCP Relay Agent Option 82 Information on page 349
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458

Copyright © 2017, Juniper Networks, Inc.

leasequery (DHCP Relay Agent)

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* ...]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Configure parameters that DHCP relay agent uses when sending DHCP leasequery messages to obtain lease information from the DHCP servers in the logical system/routing instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267

Copyright © 2017, Juniper Networks, Inc.

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

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]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation
- Configuring a DHCP Lease-Time Threshold on page 257

Copyright © 2017, Juniper Networks, Inc.

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

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*]

Release Information

Statement introduced in Junos OS Release 14.1.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring a DHCP Lease-Time Threshold on page 257

limit

Syntax

limit *max-sub-sessions*;

Hierarchy Level

[edit services service-set *services-set-name* subscriber-profile *profile-name*
max-data-sessions-per-subscriber]

Release Information

Statement introduced in Junos OS Release 11.4.

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.

Copyright © 2017, Juniper Networks, Inc.

link (Address-Assignment Pools)

Syntax link *pool-name*;

Hierarchy Level [edit access address-assignment pool *pool-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the name of a secondary address-assignment pool that is linked to a primary address-assignment pool. The secondary pool provides a backup pool for local address assignment. You can configure a chain of linked pools, but you cannot directly link more than one pool to or from any other pool.

Options ***pool-name***—Name assigned to the address-assignment pool.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585
- Configuring Address-Assignment Pool Linking on page 589

Copyright © 2017, Juniper Networks, Inc.

local (Flat-File Access Profile)

Syntax

```
local {
  flat-file-profile profile-name;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 17.1.

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.

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.

Related Documentation

- *Configuring Service Accounting in Local Flat Files*
- *Flat-File Accounting Overview*

Copyright © 2017, Juniper Networks, Inc.

local-decision (PCRF Partition)

Syntax local-decision (deny | grant) timeout *seconds*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Options

deny—Prevent subscriber logins from occurring.

grant—Allow subscriber logins to occur when the PCRF is not available or not responding.

Default: deny

seconds—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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

local-server-group (DHCP Relay Agent Option)

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 ...]

Release Information
Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

location (dhcp-relay)

Syntax location *name*;

Hierarchy Level [edit forwarding-optionsdhcp-relay relay-option-82vendor-specific],
[edit forwarding-options dhcp-relaygroup *group-name* relay-option-82vendor-specific]

Release Information Command introduced in Junos OS Release 16.2.

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. The 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.

Copyright © 2017, Juniper Networks, Inc.

location (dhcpv6)

Syntax location *name*;

Hierarchy Level [edit forwarding-optionsdhcp-relaydhcpv6 relay-option-vendor-specific],
[edit forwarding-optionsdhcp-relaygroup *group-name* dhcpv6 relay-option-vendor-specific]

Release Information Command introduced in Junos OS Release 16.2.

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. The Enterprise ID is 2636. The the location is option-data 2 (the hostname is option-data 1). The DHCPv6 relay strips the sub-option 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.

Copyright © 2017, Juniper Networks, Inc.

logical-interface-unit-range

Syntax logical-interface-unit-range (high | low)

Hierarchy Level **[edit system services extensible-subscriber-services]**

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure the unit number 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—Use upper limit of the logical interface unit range.

Range: 1 through 16,385

Default: 1

low—Use lower limit of the logical interface unit range.

Range: 1 through 16,385

Default: 16,385

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

logical-system (Diameter Peer)

Syntax logical-system *logical-system-name* <routing-instance *routing-instance-name* > ;

Hierarchy Level [edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

logical-system (Diameter Transport)

Syntax logical-system *logical-system-name* <routing-instance *routing-instance-name* >;

Hierarchy Level [edit diameter transport *transport-name*]

Release Information Statement introduced in Junos OS Release 11.2.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring the Diameter Transport on page 725

Copyright © 2017, Juniper Networks, Inc.

logical-system (Subscriber Secure Policy)

Syntax logical-system *logical-system-name*;

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 15.1R3 for enhanced subscriber management on MX Series routers.

Description 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.

Options ***logical-system-name***—Name of the logical system.

Default: Logical system **default**

Required Privilege Level flow-tap—To view this statement in the configuration.
flow-tap-control—To add this statement to the configuration.

Related Documentation

- *Subscriber Secure Policy Overview*
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 579
- *Configuring Support for Subscriber Secure Policy Mirroring*

Copyright © 2017, Juniper Networks, Inc.

logical-system-name (Static Subscribers)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

logical-system-name (DHCP Local Server)

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 ...]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

logical-system-name (DHCP Relay Agent)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

logout-response-timeout (PCRF Partition)

Syntax logout-response-timeout *seconds*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

ltv-syslog-interval (System Process)

Syntax ltv-syslog-interval *seconds*;

Hierarchy Level [edit system processes dhcp-service]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation

- Configuring a DHCP Lease-Time Threshold on page 257
- *processes*

Copyright © 2017, Juniper Networks, Inc.

mac-address (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

mac-address (DHCP Relay Agent)

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 ...]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

maintain-subscriber (Subscriber Management)

Syntax

```
maintain-subscriber {
  interface-delete;
}
```

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.1.

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. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Subscriber Binding Retention During Interface Delete Events on page 301
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302

Copyright © 2017, Juniper Networks, Inc.

managed-configuration (Dynamic Router Advertisement)

Syntax (managed-configuration | no-managed-configuration);

Hierarchy Level [edit dynamic-profiles protocols router-advertisement interface (Protocols IPv6 Neighbor Discovery) *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

map (Domain Map)

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;
  padn destination-address {
    mask destination-mask;
    metric route-metric;
  }
  strip-domain;
  strip-username (left-to-right | right-to-left);
  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]

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.

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.

Copyright © 2017, Juniper Networks, Inc.

- **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. See CLI Explorer.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Domain Map on page 164

mask (Domain Map)

Syntax
mask *destination-mask*;

Hierarchy Level
[edit access domain map *domain-map-name* padn *destination-address*]

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Configure the IP mask of the destination used in the PADN parameters for a domain map.

Options
destination-mask—Subnet mask of the destination.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring PADN Parameters for a Domain Map on page 176

Copyright © 2017, Juniper Networks, Inc.

max-advertisement-interval (Dynamic Router Advertisement)

Syntax max-advertisement-interval *seconds*;

Hierarchy Level [edit dynamic-profiles protocols router-advertisement interface (Protocols IPv6 Neighbor Discovery) *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *min-advertisement-interval (Protocols IPv6 Neighbor Discovery)*
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

max-connections (DHCP Local Server)

Syntax max-connections *number-of-connections*;

Hierarchy Level
[edit system services dhcp-local-server allow-bulk-leasequery],
[edit system services dhcp-local-server dhcpv6 allow-bulk-leasequery],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the maximum number of concurrent TCP connections allowed in the logical system/routing instance. This setting helps you manage the resources that the **jdhcp** 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.

Options *number-of-connections*—Maximum number of concurrent TCP 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.

Related Documentation
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

max-data-sessions-per-subscriber

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*]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

max-db-size (Enhanced Subscriber Management)

Syntax max-db-size *size*;

Hierarchy Level [edit system configuration-database]

Release Information Statement introduced in Junos OS Release 15.1R3 for MX Series routers with MPCs.

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 must set a **max-db-size** for the Junos configuration to allow for an allocation of the memory for enhanced subscriber management.

Juniper recommends that you set **max-db-size** to 300 MB or less for MX240, MX480, and MX960 routers, and 100 MB or less for MX5, MX10, MX40, MX80 and MX104 routers. Because you must reboot the device before the new setting takes effect, doing it as a part of the initial configuration can prevent 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: 300 MB for MX240, MX480, and MX960 routers, and 100 MB for MX5, MX10, MX40, MX80 and MX104 routers

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Junos OS Enhanced Subscriber Management Overview on page 573

Copyright © 2017, Juniper Networks, Inc.

max-empty-replies (DHCP Local Server)

Syntax max-empty-replies *number-of-replies*;

Hierarchy Level
[edit system services dhcp-local-server allow-bulk-leasequery],
[edit system services dhcp-local-server dhcpv6 allow-bulk-leasequery],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description 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.

Options ***number-of-replies***—Number of empty replies the DHCP local server sends before the connection is terminated.

Range: 1 through 100

Default: 5

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

max-outstanding-requests

Syntax max-outstanding-requests *requests*;

Hierarchy Level [edit access profile *profile-name* radius-server *server-address*],
[edit access radius-server *server-address*]

Release Information Statement introduced in Junos OS Release 11.4.

Description 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.

Options ***requests***—Maximum number of outstanding requests for this RADIUS server.

Range: 0 through 2000 outstanding requests per server

Default: 1000 outstanding requests per server

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring RADIUS Server Options for Subscriber Access on page 117
- show network-access aaa statistics on page 1736
- clear network-access aaa statistics on page 1505

Copyright © 2017, Juniper Networks, Inc.

max-outstanding-requests (Diameter Applications)

Syntax max-outstanding-requests *number*;

Hierarchy Level [edit access gx-plus global],
[edit access ocs partition *partition-name*],
[edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for OCS and PCRF partitions introduced in Junos OS Release 16.2 for MX Series routers.

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.

Related Documentation

- Configuring Gx-Plus Global Attributes on page 738
- Configuring Gx-Plus on page 736
- Configuring the OCS Partition on page 769
- Configuring the PCRF Partition on page 773
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

max-pending-accounting-stops (Access Profile)

Syntax max-pending-accounting-stops *number*;

Hierarchy Level [edit access accounting-backup-options]

Release Information Statement introduced in Junos OS Release 13.1.

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.

Related Documentation

- Configuring Back-up Options for RADIUS Accounting on page 111

max-withhold-time (Access Profile)

Syntax max-withhold-time *hold-time*;

Hierarchy Level [edit access accounting-backup-options]

Release Information Statement introduced in Junos OS Release 13.1.

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.

Related Documentation

- Configuring Back-up Options for RADIUS Accounting on page 111

Copyright © 2017, Juniper Networks, Inc.

maximum-discovery-table-entries

Syntax maximum-discovery-table-entries *entry-number*;

Hierarchy Level [edit protocols ancp],
[edit protocols ancp neighbor *ip-address*]

Release Information Statement introduced in Junos OS Release 9.5.

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.

Related Documentation
- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

maximum-helper-restart-time

Syntax maximum-helper-restart-time *seconds*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 9.4.

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.

Related Documentation
- Configuring the ANCP Agent on page 648
- Specifying How Long Processes Wait for the ANCP Agent Restart to Complete on page 654

Copyright © 2017, Juniper Networks, Inc.

maximum-lease-time

Syntax maximum-lease-time *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description 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.

Options *seconds*—Maximum number of seconds the lease can be held.

Range: 30 through 4,294,967,295 seconds

Default: 86,400 (24 hours)

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585
- DHCP Attributes for Address-Assignment Pools on page 249
- **preferred-lifetime (Address-Assignment Pools) on page 1273**
- **valid-lifetime (Address-Assignment Pools) on page 1460**

Copyright © 2017, Juniper Networks, Inc.

maximum-subscribers

Syntax maximum-subscribers *limit;*

Hierarchy Level **[edit system services extensible-subscriber-services]**

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation
- show subscribers on page 1820
- show subscribers summary on page 1846

metric (Diameter Route)

Syntax metric *route-metric*;

Hierarchy Level [edit diameter network-statement *element-name* forwarding route *dne-route-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

metric (Domain Map)

Syntax metric *route-metric*;

Hierarchy Level [edit access domain map *domain-map-name* padn *destination-address*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the route metric PADN parameter for a domain map.

Options ***route-metric***—Value 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.

Related Documentation
- Configuring PADN Parameters for a Domain Map on page 176

min-advertisement-interval (Dynamic Router Advertisement)

Syntax min-advertisement-interval *seconds*;

Hierarchy Level [edit dynamic-profiles protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *max-advertisement-interval (Protocols IPv6 Neighbor Discovery)*
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

multi-address-embedded-option-response (DHCP Local Server)

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 ...]

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 13.1 and 13.2.)

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.

Related Documentation

- Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment on page 603
- Overriding Default DHCP Local Server Configuration Settings on page 321

Copyright © 2017, Juniper Networks, Inc.

name-server

Syntax name-server [*server-names*];

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description 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.

Options ***server-names***—IP addresses of the domain name servers, listed in order of preference.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Address-Assignment Pools on page 585

nas-identifier

Syntax nas-identifier *identifier-value*;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests.

Options ***identifier-value***—String to use for authentication and accounting requests.
Range: 1 through 64 characters

Required Privilege Level admin—To view this statement in the configuration.
admin–control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

nas-port-extended-format

Syntax

```
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;
  }
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
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.

Description Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width 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.

pw-width *width*—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.

vlan-width *width*—Number of bits in the VLAN ID field.

NOTE: **The total of the widths must not exceed 32 bits, or the configuration will fail.**

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

nas-port-extended-format (Interfaces)

Syntax

```
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;
}
```

Hierarchy Level [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]

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.)

Description Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width 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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

nas-port-id-delimiter (Subscriber Management)

Syntax
nas-port-id-delimiter *delimiter-character*;

Hierarchy Level
[edit access profile *profile-name* radius options]

Release Information
Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description
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.

Default
The hash (**#**) character.

Options
delimiter-character—Character used for the delimiter.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring a NAS-Port-ID with Additional Options on page 63

Copyright © 2017, Juniper Networks, Inc.

nas-port-id-format (Subscriber Management)

Syntax

```
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;
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Options **interface-text-description**, **order**, and **postpend-vlan-tags** introduced in Junos OS Release 15.1.

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**

Copyright © 2017, Juniper Networks, Inc.

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

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.

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.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring a NAS-Port-ID with Additional Options on page 63

Copyright © 2017, Juniper Networks, Inc.

nas-port-options (RADIUS Options)

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]

Release Information Statement introduced in Junos OS Release 12.3.

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. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74

Copyright © 2017, Juniper Networks, Inc.

- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

 Copyright © 2017, Juniper Networks, Inc.

nas-port-type (Subscriber Management)

Syntax

```
nas-port-type {
  ethernet {
    port-type;
  }
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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

Copyright © 2017, Juniper Networks, Inc.

- **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.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring RADIUS Server Parameters for Subscriber Access on page 116

Copyright © 2017, Juniper Networks, Inc.

nas-port-type (RADIUS Options)

Syntax nas-port-type *port-type*;

Hierarchy Level [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]

Release Information Statement introduced in Junos OS Release 12.3.

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

Copyright © 2017, Juniper Networks, Inc.

- **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.

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

Copyright © 2017, Juniper Networks, Inc.

nasreq (Diameter Application)

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Diameter Network Access Server Application (NASREQ) on page 779
- Configuring the Diameter Network Access Server Application (NASREQ) on page 781

neighbor (Define ANCP)

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]

Release Information Statement introduced in Junos OS Release 9.4.

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. See CLI Explorer.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the ANCP Agent on page 648
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

neighbor-discovery-router-advertisement (Address-Assignment Pools)

Syntax neighbor-discovery-router-advertisement *ndra-pool-name*;

Hierarchy Level [edit access address-assignment]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the name of the address-assignment pool used to assign the router advertisement prefix.

Options ***ndra-pool-name***—Name of the address-assignment pool.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring an Address-Assignment Pool Used for Router Advertisements on page 486

netbios-node-type

Syntax netbios-node-type *node-type*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the NetBIOS node type. This is equivalent to DHCP option 46.

Options ***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

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

network

Syntax network *ip-prefix</prefix-length>*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet]

Release Information Statement introduced in Junos OS Release 9.0.

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.

Related Documentation

- Configuring Address-Assignment Pools on page 585

 Copyright © 2017, Juniper Networks, Inc.

network-element (Diameter Base Protocol)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

network-services

Syntax network-services (ethernet | enhanced-ethernet | ip | enhanced-ip | lan);

Hierarchy Level [edit chassis]

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.

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 Trio 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 Trio 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, MX 2010, and 2020 3D Universal Edge Routers.

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.

Related Documentation
- *Network Services Mode Overview*
- *Firewall Filters and Enhanced Network Services Mode Overview* in the *Junos OS Broadband Subscriber Management and Services Library*

Copyright © 2017, Juniper Networks, Inc.

- *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers*
- *Configuring Enhanced IP Network Services for a Virtual Chassis*

Copyright © 2017, Juniper Networks, Inc.

no-bind-on-request (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 10.4.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324
- Disabling Automatic Binding of Stray DHCP Requests on page 348

no-unsolicited-ra (Enhanced Subscriber Management)

Syntax no-unsolicited-ra;

Hierarchy Level [edit system services subscriber-management overrides]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575

Copyright © 2017, Juniper Networks, Inc.

no-vlan-interface-name

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]

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.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Including a Textual Description in DHCP Options on page 354
- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring DHCPv6 Relay Agent Options on page 456

Copyright © 2017, Juniper Networks, Inc.

ocs (Diameter Applications)

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;
    }
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 16.2.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

on-demand-ip-address

Syntax on-demand-ip-address;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options], [edit interfaces pp0 unit *unit-number* ppp-options], [edit protocols ppp-service]

Release Information Statement introduced in Junos OS Release 13.1.

Description For IPv4 and IPv6 dual-stack PPP subscribers, enables on-demand allocation and de-allocation of an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address.

Configuration changes take effect as follows:

- When you change this setting for a dynamic PPP interface (at the **[edit dynamic-profiles]** hierarchy level), the change takes effect only for new subscriber logins.
- When you change this setting for a static PPP interface (at the **[edit interfaces pp0]** hierarchy level, the subscribers on the interface are logged out.
- When you change this setting globally (at the **[edit protocols ppp-service]** hierarchy level), the change takes effect only for new subscriber logins.

If you enable on-demand allocation at both the interface and global levels, the global configuration takes precedence and changes take effect for new subscriber logins.

Default This functionality is disabled by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation on page 461
- Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 469
- Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470
- Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers on page 470

Copyright © 2017, Juniper Networks, Inc.

on-link (Dynamic Router Advertisement)

Syntax (on-link | no-on-link);

Hierarchy Level [edit dynamic-profiles *profile-name* protocols router-advertisement interface *interface-name*prefix *prefix*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

option

Syntax

```
option {
  [ (id-number option-type option-value)
  (id-number array option-type option-value) ];
}
```

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.
hex-string option type introduced in Junos OS Release 13.3.

Description Specify user-defined options that are added to client packets.

Options

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).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

option-60 (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

option-60 (DHCP Relay Agent)

Syntax option-60;

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

 Copyright © 2017, Juniper Networks, Inc.

option-82 (DHCP Relay Agent)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

option-82 (DHCP Local Server Authentication)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

option-82 (DHCP Local Server Pool Matching)

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]

Release Information
Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- Extended DHCP Local Server Overview on page 234
- Address-Assignment Pools Overview on page 583

Copyright © 2017, Juniper Networks, Inc.

option-82 (Address-Assignment Pools)

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]

Release Information Statement introduced in Junos OS Release 9.0.

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.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

option-match

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*]

Release Information Statement introduced in Junos OS Release 9.0.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

option-number (DHCP Relay Agent Option)

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 ...]

Release Information
Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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.

Related Documentation
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

options (Access Profile)

Syntax

```
options {
    accounting-session-id-format (decimal | description);
    calling-station-id-delimiter delimiter-character;
    calling-station-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    chap-challenge-in-request-authenticator;
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm (direct | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    access-loop-id-local;
    interface-description-format {
        exclude-adapter;
        exclude-sub-interface;
    }
    ip-address-change-notify message;
    juniper-dsl-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
            port-width width;
            pw-width width;
            slot-width width;
            vci-width width;
            vpi-width width;
        }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        order {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            postpend-vlan-tags;
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
    postpend-vlan-tags;
  }
  nas-port-type {
    ethernet {
      port-type;
    }
  }
  revert-interval interval;
  service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
  }
  vlan-nas-port-stacked-format;
}
```

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the options used by RADIUS authentication and accounting servers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- RADIUS Server Options for Subscriber Access on page 4

Copyright © 2017, Juniper Networks, Inc.

order

Syntax order [*accounting-method*];

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

origin (Diameter Base Protocol)

Syntax origin realm *realm-name* host *hostname*

Hierarchy Level [edit diameter]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring the Origin Attributes of the Diameter Instance on page 724

Copyright © 2017, Juniper Networks, Inc.

other-bytes

Syntax other-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 16.1.

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 lines (GPON)—as DSL type OTHER.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

other-overhead-adjust

Syntax other-overhead-adjust *percentage*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 16.1.

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 lines (GPON)—as DSL type OTHER.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

other-stateful-configuration (Dynamic Router Advertisement)

Syntax (other-stateful-configuration | no-other-stateful-configuration);

Hierarchy Level [edit dynamic-profiles protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation

- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399

Copyright © 2017, Juniper Networks, Inc.

overhead-accounting (ANCP)

Syntax overhead-accounting;

Hierarchy Level [edit protocols ancp interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation

- Configuring the ANCP Agent on page 648

 Copyright © 2017, Juniper Networks, Inc.

overrides (DHCP Local Server)

Syntax

```
overrides {
  allow-no-end-option;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match <option60-and-option82 | incoming-interface>;
  client-negotiation-match incoming-interface;
  delegated-pool;
  delete-binding-on-renegotiation;
  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 ...]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Support for the **allow-no-end** option introduced in Junos OS Release 14.1X53-D15 for EX Series switches.
Statement introduced in Junos OS Release 12.3X48-D10 for SRX Series devices.

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.

Copyright © 2017, Juniper Networks, Inc.

- 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. See CLI Explorer.

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.

Related Documentation

- Extended DHCP Local Server Overview on page 234
- Overriding Default DHCP Local Server Configuration Settings on page 321
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327
- *Configuring a DHCP Server on Switches (CLI Procedure)*

Copyright © 2017, Juniper Networks, Inc.

overrides (DHCP Relay Agent)

Syntax

```
overrides {
  allow-no-end-option;
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match <option60-and-option82 | incoming-interface>;
  client-negotiation-match incoming-interface;
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*],
[edit forwarding-options dhcp-relay group *group-name*],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information

Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Override the default configuration settings for the extended DHCP relay agent. Specifying the **overrides** statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

The following statements are supported at both the **[edit ... dhcp-relay]** and **[edit ... dhcpv6]** hierarchy levels.

- **allow-snooped-clients**

Copyright © 2017, Juniper Networks, Inc.

- **asymmetric-lease-time**
- **delete-binding-on-renegotiation**
- **dual-stack**
- **interface-client-limit**
- **no-allow-snooped-clients**
- **no-bind-on-request**
- **relay-source**
- **send-release-on-delete**

The following statements are supported at the **[edit ... dhcpv6]** hierarchy levels only.

- **asymmetric-prefix-lease-time**

All other statements are supported at the **[edit ... dhcp-relay]** hierarchy levels only.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 327

Copyright © 2017, Juniper Networks, Inc.

overrides (Enhanced Subscriber Management)

Syntax

```
overrides {
  no-unsolicited-ra;
  shmlog {
    disable;
    file filename <files maximum-no-files> <size maximum-file-size>;
    filtering enable;
    log-name {
      all;
      logname {
        <brief | detail | extensive | none | terse>;
        <file-logging |no-file-logging>;
      }
      }
    log-type (debug | info | notice);
  |
}
```

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Override the default configuration settings for the Junos OS enhanced subscriber management software for subscriber management.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575

Copyright © 2017, Juniper Networks, Inc.

override-password (Domain Map)

Syntax override-password *password*

Hierarchy Level [edit access domain *domain-name*map]

Release Information Statement introduced in Junos OS 14.1.

Description Specify a password to be used as the authenticating password for all outgoing authentication requests for subscribers who match the domain map.

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.

Related Documentation

- strip-username on page 1387

padn (Domain Map)

Syntax

```
padn destination-address {
  mask destination-mask;
  metric route-metric;
}
```

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure PADN parameters for a domain map.

Options ***destination-address***—IP address of the destination.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring PADN Parameters for a Domain Map on page 176

Copyright © 2017, Juniper Networks, Inc.

parse-direction (Domain Map)

Syntax parse-direction (left-to-right | right-to-left);

Hierarchy Level [edit access domain]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation

- Specifying the Parsing Direction for Domain and Realm Names on page 173
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

parse-order (Domain Map)

Syntax parse-order (domain-first | realm-first);

Hierarchy Level [edit access domain]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation

- Specifying the Parsing Direction for Domain and Realm Names on page 173
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

partition

Syntax

```
partition partition-name {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm;
}
```

Hierarchy Level [edit jsrc]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a JSRC partition.

Options *partition-name*—Name of the JSRC partition.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring JSRC on page 787
- Configuring the JSRC Partition on page 788

Copyright © 2017, Juniper Networks, Inc.

partition (Gx-Plus)

Syntax

```
partition partition-name {
  diameter-instance instance-name;
  destination-host hostname;
  destination-realm realm;
}
```

Hierarchy Level [edit access gx-plus]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure a Gx-Plus partition.

Options *partition-name*—Name of the Gx-Plus partition.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Gx-Plus on page 736
- Configuring the Gx-Plus Partition on page 737

Copyright © 2017, Juniper Networks, Inc.

partition (NASREQ Diameter Application)

Syntax

```
partition partition-name {
  destination-host hostname;
  destination-realm realm-name;
  diameter-instance master;
}
```

Hierarchy Level [edit access nasreq]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Diameter Network Access Server Application (NASREQ) on page 779
- Configuring the Diameter Network Access Server Application (NASREQ) on page 781

Copyright © 2017, Juniper Networks, Inc.

partition (OCS)

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]

Release Information Statement introduced in Junos OS Release 16.2.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

partition (PCRF)

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) 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;
  }
  subscription-id-type number;
  update-response-timeout seconds;
}
```

Hierarchy Level [edit access pcrf]

Release Information Statement introduced in Junos OS Release 16.2.

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.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

password (Static Subscribers)

Syntax

```
password password-string;
  username-include {
    domain-name domain-name;
    username-include;
    logical-system-name;
    routing-instance-name;
    user-prefix user-prefix-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]

Release Information

Statement introduced in Junos OS Release 9.6.

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.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system-level—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Authentication Password on page 807
- Configuring the Static Subscriber Group Authentication Password on page 811

Copyright © 2017, Juniper Networks, Inc.

password (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

password (DHCP Relay Agent)

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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281
- Configuring Passwords for Usernames on page 285

Copyright © 2017, Juniper Networks, Inc.

pcrf (Diameter Applications)

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) 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;
    }
    subscription-id-type number;
    update-response-timeout seconds;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 16.2.

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. See CLI Explorer.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

peer (Diameter Base Protocol)

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]

Release Information
Statement introduced in Junos OS Release 9.6.

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. See CLI Explorer.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

peer (Diameter Network Element)

Syntax

```
peer peer-name {
    priority priority-value;
}
```

Hierarchy Level [edit diameter network-element *element-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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 statement is explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

peer-ip-address-optional

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]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- IPCP Negotiation with Optional Peer IP Address on page 468
- *Dynamic Profiles Overview*

 Copyright © 2017, Juniper Networks, Inc.

peer-origin (Diameter Peer)

Syntax peer-origin realm *realm-name* host *hostname*;

Hierarchy Level [edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring Diameter Peers on page 724
- Configuring Diameter on page 723

Copyright © 2017, Juniper Networks, Inc.

pool (Address-Assignment Pools)

Syntax

```
pool pool-name {
  active-drain;
  family family {
    dhcp-attributes {
      [ protocol-specific attributes ]
    }
    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 address-assignment]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

pool (DHCP Local Server Overrides)

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]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation
- Enabling Processing of Client Information Requests on page 278
- Overriding Default DHCP Local Server Configuration Settings on page 321

pool-match-order

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]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1.

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. See CLI Explorer.

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.

Related Documentation
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 252
- Extended DHCP Local Server Overview on page 234

Copyright © 2017, Juniper Networks, Inc.

port

Syntax port *port-number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the port number on which to contact the RADIUS server.

Options ***port-number***—Port number on which to contact the RADIUS server.
Default: 1812 (as specified in RFC 2865)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

port (Diameter Peer)

Syntax port *port-number*;

Hierarchy Level [edit diameter peer *peer-name* connect-actively]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

pre-ietf-mode

Syntax pre-ietf-mode

Hierarchy Level [edit protocols ancp],
[edit protocols ancp neighbor *ip-address*]

Release Information Statement introduced in Junos OS Release 9.4.

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.

Related Documentation

- Configuring the ANCP Agent on page 648
- Configuring the ANCP Agent for Backward Compatibility on page 653
- Configuring ANCP Neighbors on page 649

Copyright © 2017, Juniper Networks, Inc.

preauthentication-order (Access Profile)

Syntax preauthentication-order [*preauthentication-method*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 13.3.

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.

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

Copyright © 2017, Juniper Networks, Inc.

preauthentication-port

Syntax preauthentication-port *port-number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description 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.

Options ***port-number***—Port number used for preauthentication requests to contact the RADIUS server.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- RADIUS Attributes for LLID Preauthentication Requests on page 150

Copyright © 2017, Juniper Networks, Inc.

preauthentication-secret

Syntax preauthentication-secret *password*;

Hierarchy Level [edit access profile *profile-name* radius-server *server-address*],
[edit access radius-server *server-address*]

Release Information Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description Configure the password to use with the RADIUS server for logical line identification (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.

Options ***password***—Password to use. To include spaces enclose the character string in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- RADIUS Attributes for LLID Preauthentication Requests on page 150

Copyright © 2017, Juniper Networks, Inc.

preauthentication-server (Access Profile)

Syntax preauthentication-server *ip-address*;

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 13.3.

Description 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.

Options ***ip-address***—IPv4 address.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- RADIUS Logical Line Identifier (LLID) Overview on page 149
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

Copyright © 2017, Juniper Networks, Inc.

preferred-lifetime (Address-Assignment Pools)

Syntax preferred-lifetime *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 13.3.

Description 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.

Options *seconds*—Number of seconds that the IPv6 prefix is active.

Range: 30 through 4,294,967,295 seconds

Default: 86,400 (24 hours)

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585
- DHCP Attributes for Address-Assignment Pools on page 249
- maximum-lease-time on page 1190
- valid-lifetime (Address-Assignment Pools) on page 1460

Copyright © 2017, Juniper Networks, Inc.

preferred-lifetime (Dynamic Router Advertisement)

Syntax preferred-lifetime *seconds*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols router-advertisement interface *interface-name*prefix *prefix*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation

- *valid-lifetime*
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

prefix (DHCP Relay Agent)

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)]

Release Information
Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
Support at the **[edit ... relay-agent-remote-id]** and **[edit ... remote-id]** hierarchy levels introduced in Junos OS Release 14.1.

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.

Related Documentation
- Including a Prefix in DHCP Options on page 352
- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring DHCPv6 Relay Agent Options on page 456

Copyright © 2017, Juniper Networks, Inc.

prefix (Address-Assignment Pools)

Syntax prefix *ipv6-prefix*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

prefix (Dynamic Router Advertisement)

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*]

Release Information Statement introduced in Junos OS Release 10.1.

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. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

priority (Diameter Peer)

Syntax priority *priority-value*;

Hierarchy Level [edit diameter network-element *element-name* peer *peer-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

Copyright © 2017, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    address-change-immediate-update
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    duplication-vrf {
      access-profile-name profile-name;
      vrf-name vrf-name;
    }
    immediate-update;
    order [ accounting-method ];
    send-acct-status-on-config-change;
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
  }
  accounting-order (radius | [accounting-order-data-list]);
  authentication-order [ authentication-methods ];
  client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
      allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
      }
      pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
      ike-policy policy-name;
      interface-id string-value;
    }
    l2tp {
      aaa-access-profile profile-name;
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions number;
      maximum-sessions-per-tunnel number;
      multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
      }
      override-result-code session-out-of-resource;
      ppp-authentication (chap | pap);
      ppp-profile profile-name;
      sessions-limit-group limit-group-name;
      shared-secret shared-secret;
    }
    pap-password pap-password;
```


Copyright © 2017, Juniper Networks, Inc.

```
    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 {
        ...
      }
      ignore {
        framed-ip-netmask;
        input-filter;
        logical-system:routing-instance;
        output-filter;
      }
    }
    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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
            exclude-adapter;
            exclude-sub-interface;
          }
          juniper-dsl-attributes;
          nas-identifier identifier-value;
          nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            pw-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
              adapter-width width;
              port-width width:
              slot-width width;
              vci-width width:
              vpi-width width;
            }
          }
          nas-port-id-delimiter delimiter-character;
          nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            order {
              agent-circuit-id;
              agent-remote-id;
              interface-description;
              interface-text-description;
              nas-identifier;
              postpend-vlan-tags;
            }
            postpend-vlan-tags;
          }
          nas-port-type {
            ethernet {
              port-type;
            }
          }
          revert-interval interval;
          service-activation {
            dynamic-profile (optional-at-login | required-at-login);
            extensible-service (optional-at-login | required-at-login);
          }
          vlan-nas-port-stacked-format;
        }
        preauthentication-server ip-address;
      }
      radius-server server-address {
        accounting-port port-number;
        accounting-retry number;
        accounting-timeout seconds;
```


Copyright © 2017, Juniper Networks, Inc.

```
    dynamic-request-port
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    max-outstanding-requests value;
    source-address source-address;
    timeout seconds;
  }
  service {
    accounting {
      statistics (time | volume-time);
      update-interval minutes;
    }
    accounting-order (activation-protocol | local | radius);
  }
  session-options {
    client-idle-timeout minutes;
    client-idle-timeout-ingress-only;
    client-session-timeoutminutes;
    strip-user-name {
      delimiter [ delimiter ];
      parse-direction (left-to-right | right-to-left);
    }
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure PPP CHAP, or a profile and its subscriber access, L2TP, or PPP properties.

Options ***profile-name***—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- *Configuring the PPP Authentication Protocol*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring an L2TP LNS with Inline Service Interfaces*
- *Configuring PPP Properties for a Client-Specific Profile*
- Configuring Service Accounting with JSRC on page 794
- *Configuring Service Accounting in Local Flat Files*
- AAA Service Framework Overview on page 3
- **show network-access aaa statistics on page 1736**
- **clear network-access aaa statistics on page 1505**

Copyright © 2017, Juniper Networks, Inc.

process-inform

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]

Release Information

Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Copyright © 2017, Juniper Networks, Inc.

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 statement is explained separately. See CLI Explorer.

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.

Related Documentation

- Enabling Processing of Client Information Requests on page 278
- Overriding Default DHCP Local Server Configuration Settings on page 321

Copyright © 2017, Juniper Networks, Inc.

protocols (Dynamic Profiles)

Syntax

```
protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-policy;
      immediate-leave
      no-accounting;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group group {
          source source;
        }
      }
      version version;
    }
    mld {
      interfaceinterface-name {
        disable;
        (accounting | no-accounting);
        group-policy;
        immediate-leave;
        oif-map;
        passive;
        ssm-map ssm-map-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;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefix prefix;
        reachable-time milliseconds;
        retransmit-timer milliseconds;
      }
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name*]

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.

Description Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.

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).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring IGMP*
- *Examples: Configuring MLD*

Copyright © 2017, Juniper Networks, Inc.

provisioning-order Diameter Applications)

Syntax provisioning-order (gx-plus | jsrc | pcrf);

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Support for Gx-Plus introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
pcrf option added in Junos OS Release 16.2.

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.

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.

Related Documentation

- Configuring JSRC on page 787
- Provisioning Subscribers with JSRC on page 790
- Configuring Gx-Plus on page 736
- Provisioning Subscribers with Gx-Plus on page 739
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

proxy-mode

Syntax proxy-mode;

Hierarchy Level [edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.

You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- DHCP Relay Proxy Overview on page 243
- Extended DHCP Relay Agent Overview on page 240
- Enabling DHCP Relay Proxy Mode on page 369

Copyright © 2017, Juniper Networks, Inc.

qos-adjust

Syntax

```
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;
}
```

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify that the ANCP agent reports data rates for downstream traffic to CoS. Configure the values by which the actual downstream data rates are adjusted. When this statement is not configured, the ANCP agent does not report traffic rates to CoS.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-adsl

Syntax qos-adjust-adsl *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-adsl2

Syntax qos-adjust-adsl2 *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-adsl2-plus

Syntax qos-adjust-adsl2-plus *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-other

Syntax qos-adjust-other *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 16.1.

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 lines (GPON)—as DSL type OTHER.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-sdsl

Syntax qos-adjust-sdsl *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-vdsl

Syntax qos-adjust-vdsl *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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 VDSL line. The ANCP agent reports the adjusted rate only to AAA.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

qos-adjust-vdsl2

Syntax qos-adjust-vdsl2 *adjustment-factor*;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 11.4.

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 VDSL2 line. The ANCP agent reports the adjusted rate only to AAA.

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.

Related Documentation

- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 685
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

radius (Access Profile)

Syntax

```
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude
      ...
    }
    ignore {
      framed-ip-netmask;
      input-filter;
      logical-system-routing-instance;
      output-filter;
    }
  }
  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-sub-interface;
    }
    ip-address-change-notify message;
    juniper-dsl-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 {
```

Copyright © 2017, Juniper Networks, Inc.

```
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        order {
          agent-circuit-id;
          agent-remote-id;
          interface-description;
          interface-text-description;
          nas-identifier;
          postpend-vlan-tags;
        }
        postpend-vlan-tags;
      }
      nas-port-type {
        ethernet {
          port-type;
        }
      }
      revert-interval interval;
      service-activation {
        dynamic-profile (optional-at-login | required-at-login);
        extensible-service (optional-at-login | required-at-login);
      }
      vlan-nas-port-stacked-format;
    }
    preauthentication-server ip-address;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- RADIUS Server Options for Subscriber Access on page 4

Copyright © 2017, Juniper Networks, Inc.

radius-disconnect (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

Copyright © 2017, Juniper Networks, Inc.

- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 312

 Copyright © 2017, Juniper Networks, Inc.

radius-flow-tap

Syntax

```
radius-flow-tap {
  forwarding-class class-name;
  interfaces interface-name;
  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;
        }
      }
    }
  }
  routing-instance routing-instance-name;
  source-ipv4-address ipv4-address;
)
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 9.4.

Description Assign parameters that are used with subscriber secure policy mirroring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level flow-tap—To view this statement in the configuration.
flow-tap-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- *Subscriber Secure Policy Overview*
- *Configuring Support for Subscriber Secure Policy Mirroring*

radius-options (Edit Access)

Syntax

```
radius-options {
  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]

Release Information Statement introduced in Junos OS Release 8.3.

Description Configure RADIUS options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable on page 121
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

radius-options (Interfaces)

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*]

Release Information Statement introduced in Junos OS Release 12.3.

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. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

Copyright © 2017, Juniper Networks, Inc.

radius-server

Syntax

```
radius-server server-address {
  accounting-port port-number;
  accounting-retry number;
  accounting-timeout seconds;
  dynamic-request-port
  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*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
dynamic-request-port option added in Junos OS Release 14.2 for MX Series routers.
preauthentication-port and **preauthentication-secret** options added in Junos OS Release 15.1 for MX Series routers.
Support for IPv6 ***server-address*** introduced in Junos OS Release 16.1.

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.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring RADIUS Authentication for L2TP*
- *Configuring the PPP Authentication Protocol*
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

- show network-access aaa statistics on page 1736
- clear network-access aaa statistics on page 1505

range (Address-Assignment Pools)

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)]

Release Information Statement introduced in Junos OS Release 9.0.
IPv6 support introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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.

Related Documentation

- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

rapid-commit (DHCPv6 Local Server)

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 ...]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation

- Enabling DHCPv6 Rapid Commit Support on page 449
- Overriding Default DHCP Local Server Configuration Settings on page 321

Copyright © 2017, Juniper Networks, Inc.

reachable-time (Dynamic Router Advertisement)

Syntax reachable-time *milliseconds*;

Hierarchy Level [edit protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

realm-delimiter (Domain Map)

Syntax realm-delimiter [*delimiter-character*];

Hierarchy Level [edit access domain]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation
- Specifying Domain and Realm Name Delimiters on page 172
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

realm-parse-direction (Domain Map)

Syntax realm-parse-direction (left-to-right | right-to-left);

Hierarchy Level [edit access domain]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation

- Specifying the Parsing Direction for Domain and Realm Names on page 173
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

Copyright © 2017, Juniper Networks, Inc.

reconfigure (DHCP Local Server)

Syntax

```
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  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*]

Release Information

Statement introduced in Junos OS Release 10.0.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

relay-agent-interface-id (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-agent-interface-id (DHCPv6 Relay Agent)

Syntax

```
relay-agent-interface-id {
  include-l2-interface-for-irb;
  keep-incoming-interface-id;
  prefix prefix;
  use-interface-description (logical | device);
  use-option-82;
}
```

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* routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 ...]

Release Information

Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1020
- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458

Copyright © 2017, Juniper Networks, Inc.

relay-agent-interface-id (DHCPv6 Relay Agent Username)

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]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation
- DHCPv6 Relay Agent Overview on page 455
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-agent-remote-id (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-agent-remote-id (DHCPv6 Relay Agent)

Syntax

```
relay-agent-remote-id {
  keep-incoming-remote-id;
  prefix prefix;
  use-interface-description (logical | device);
  use-option-82 <strict>;
}
```

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 ...]

Release Information

Statement introduced in Junos OS Release 14.1.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Extended DHCP Relay Agent Overview on page 240
- DHCPv6 Relay Agent Overview on page 455
- DHCPv6 Relay Agent Options on page 456
- Configuring DHCPv6 Relay Agent Options on page 456

Copyright © 2017, Juniper Networks, Inc.

relay-agent-remote-id (DHCPv6 Relay Agent Username)

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]

Release Information
Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Related Documentation
- DHCPv6 Relay Agent Overview on page 455
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-agent-subscriber-id (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-agent-subscriber-id (DHCPv6 Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation
- DHCPv6 Relay Agent Overview on page 455
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

relay-option (DHCP Relay Agent)

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 ...]

Release Information

Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

relay-option-vendor-specific (dhcpv6)

Syntax

```
relay-option-vendor-specific {
  host-name;
  location;
}
```

Hierarchy Level [edit forwarding-optionsdhcp-relaydhcpv6],

Release Information Command introduced in Junos OS Release 16.2.

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.

Copyright © 2017, Juniper Networks, Inc.

relay-option-82

Syntax

```
relay-option-82 {
  circuit-id {
    include-irb-and-l2;
    keep-incoming-circuit-id;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
  }
  remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
  }
  server-id-override
  vendor-specific{
    host-name;
    location;
  }
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay group *group-name*],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name*],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

To enable insertion of option 82 information in DHCP packets, you must specify at least one of the **circuit-id** or **remote-id** statements.

You can use the **relay-option-82** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level to control insertion of option 82 information globally, or at the **[edit forwarding-options dhcp-relay** group ***group-name*****]** hierarchy level to control insertion of option 82 information for a named group of interfaces.

Copyright © 2017, Juniper Networks, Inc.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the **delete relay-option-82** statement.

The remaining statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using DHCP Relay Agent Option 82 Information on page 349
- dhcp-relay on page 1020

relay-server-group (DHCP Relay Agent Option)

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 ...]

Release Information
Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

relay-source

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation
- Configuring DHCP Relay Agent Source Address on page 360

Copyright © 2017, Juniper Networks, Inc.

remote-id

Syntax remote-id *value* range *named-range*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes option-match option-82],
[edit access protocol-attributes *attribute-set-name* option-match option-82]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the address-assignment pool named range to use based on the particular option 82 Agent Remote ID value.

Options 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.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

remote-id (DHCP Relay Agent)

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]

Release Information

Statement introduced in Junos OS Release 14.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring Option 82 Information on page 350

Copyright © 2017, Juniper Networks, Inc.

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

Syntax remote-id-mismatch *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*]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify that the router match the remote ID string for DHCPv4 and DHCPv6 clients. The router then inspects incoming Renew and Rebind messages and compares the Remote ID in the message against the initial value that DHCP stored when creating the client entry. If there is a mismatch between the initial value and the Renew or Rebind message value, DHCP local server performs the specified action.

Options ***action***—Action the DHCP local server takes when it detects a remote ID mismatch.

- **disconnect**—DHCP local server disconnects the client, which forces the client to reconnect.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring DHCP-Initiated Service Change Based on Remote ID on page 366

 Copyright © 2017, Juniper Networks, Inc.

replace-ip-source-with

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]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Extended DHCP Relay Agent Overview on page 240
- Replacing the DHCP Relay Request and Release Packet Source Address on page 337

Copyright © 2017, Juniper Networks, Inc.

report-interface-descriptions (Edit Access)

Syntax report-interface-descriptions;

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 13.1.

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.

Related Documentation

- exclude on page 1075
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview on page 8

Copyright © 2017, Juniper Networks, Inc.

report-local-rule (PCRF Partition)

Syntax report-local-rule;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

report-resource-allocation (PCRF Partition)

Syntax report-resource-allocation;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744

Copyright © 2017, Juniper Networks, Inc.

- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

report-successful-resource-allocation (PCRF Partition)

Syntax report-successful-resource-allocation;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

- Understanding Upstream and Downstream Messages for the PCRF on page 766

request-max-tcp-connections (System Process)

Syntax request-max-tcp-connections *max-tcp-connections*;

Hierarchy Level [edit system processes dhcp-service]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

request network-access aaa replay pending-accounting-stops

Syntax request network-access aaa replay pending-accounting-stops

Release Information Command introduced in Junos OS Release 13.1.

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

Related Documentation
- Forcing the Router to Contact the Accounting Server Immediately on page 115
- show accounting pending-accounting-stops on page 1548

List of Sample Output request network-access aaa replay pending-accounting-stops on page 1334

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
```


Copyright © 2017, Juniper Networks, Inc.

request network-access aaa subscriber add session-id

Syntax request network-access aaa subscriber add session-id *subscriber-session-id* service-profile *profile-name*

Release Information Command introduced in Junos OS Release 11.2.

Description Locally activate a dynamic subscriber service for a subscriber who is currently logged in to the network.

Options ***profile-name***—Name of service-profile to activate.

subscriber-session-id—ID of the subscriber session for which the service will be added.

Required Privilege Level view

Related Documentation

- CLI-Activated Subscriber Services on page 609
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- request network-access aaa subscriber delete session-id on page 1337

List of Sample Output request network-access aaa subscriber add session-id service-profile on page 1336

Output Fields When you enter this command, you are provided feedback on the status of your request. Table 72 on page 1335 lists possible error messages that might be returned if the service activation fails.

Table 72: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Command failed: *reason*	–	–
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.
Provisioning is already active	Remote provisioning by a JSRC server or Gx-plus server is active.	–
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.

Copyright © 2017, Juniper Networks, Inc.

Table 72: 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 add session-id service-profile

```
user@host> request network-access aaa subscriber add session-id 49 service-profile
service-bronze
Successful completion
```


Copyright © 2017, Juniper Networks, Inc.

request network-access aaa subscriber delete session-id

Syntax request network-access aaa subscriber delete session-id *subscriber-session-id* service-profile *profile-name*

Release Information Command introduced in Junos OS Release 11.2.

Description Deactivate a dynamic subscriber service for a subscriber who is currently logged in to the network.

Options ***profile-name***—Name of the service-profile to deactivate. To deactivate a single instance of a subscriber service that has multiple instances, you can specify the service-profile name and its service parameters.

subscriber-session-id—ID of the subscriber session for which the service will be deleted.

Required Privilege Level view

Related Documentation

- CLI-Activated Subscriber Services on page 609
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622
- **request network-access aaa subscriber add session-id on page 1335**

List of Sample Output **request network-access aaa subscriber delete session-id service-profile on page 1338**
request network-access aaa subscriber delete session-id service-profile (Deactivating a Single Server Instance) on page 1338
request network-access aaa subscriber delete session-id service-profile (Deactivating All Server Instances) on page 1338

Output Fields When you enter this command, you are provided feedback on the status of your request. Table 73 on page 1337 lists possible error messages that might be returned if the service deactivation fails.

Table 73: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Command failed: *reason*	Error condition that caused the command to fail.	Correct the error condition.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 73: Service Activation/Deactivation Error Messages *(continued)*

Message	Description	Corrective Action
Provisioning is already active	Remote provisioning by a JSRC server or Gx-plus server is active.	Disable provisioning.
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.
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 delete session-id service-profile

```
user@host> request network-access aaa subscriber delete session-id 49 service-profile
service-silver
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating a Single Server Instance)

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile
economy-service(up-filter,down-filter)
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating All Server Instances)

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile
economy-service
Successful completion
```


Copyright © 2017, Juniper Networks, Inc.

request network-access aaa subscriber modify session-id

Syntax request network-access aaa subscriber modify session-id *subscriber-session-id predefined-variable variable-option*

Release Information Command introduced in Junos OS Release 11.2.

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

Related Documentation

- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 613
- CLI-Activated Subscriber Services on page 609

List of Sample Output request network-access aaa subscriber modify session-id on page 1339

Output Fields When you enter this command, you are provided feedback on the status of your request. Table 74 on page 1339 lists possible messages that might be returned.

Table 74: 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
```

Copyright © 2017, Juniper Networks, Inc.

request-rate

Syntax request-rate *rate*;

Hierarchy Level [edit access radius-options]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation

- Global RADIUS Options for Subscriber Access on page 8
- Global RADIUS Options for Subscriber Access on page 8
- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Router or Switch Interaction with RADIUS Servers on page 113

Copyright © 2017, Juniper Networks, Inc.

requested-ip-network-match (DHCP Local Server)

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 ...]

Release Information Statement introduced in Junos OS Release 12.2.

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 15
- 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.

Related Documentation

- Specifying the Subnet for DHCP Client Address Assignment on page 255

Copyright © 2017, Juniper Networks, Inc.

restricted-requestor (DHCP Local Server)

Syntax restricted-requestor;

Hierarchy Level
[edit system services dhcp-local-server allow-bulk-leasequery],
[edit system services dhcp-local-server allow-leasequery],
[edit system services dhcp-local-server dhcpv6 allow-bulk-leasequery],
[edit system services dhcp-local-server dhcpv6 allow-leasequery],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify that the DHCP local server responds to a leasequery or 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 leasequery and bulk leasequery requests, the giaddr of the requestor must match the giaddr of the client.
- For DHCPv6 leasequery and bulk 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.

Related Documentation
- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

retransmit-timer (Dynamic Router Advertisement)

Syntax retransmit-timer *milliseconds*;

Hierarchy Level [edit protocols router-advertisement interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

Copyright © 2017, Juniper Networks, Inc.

retry

Syntax retry *attempts*;

Hierarchy Level [edit access radius-server *server-address*];
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the number of times that the router or switch 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.

Options ***attempts***—Number of times that the router is allowed to attempt to contact a RADIUS server.

Range: 1 through 100

Default: 3

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*
- timeout on page 1402

Copyright © 2017, Juniper Networks, Inc.

revert-interval

Syntax revert-interval *interval*;

Hierarchy Level [edit access profile *profile-name* radius options];
[edit access radius-options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the 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.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

route (Diameter Network Element)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Network Elements on page 726

router (Address-Assignment Pools)

Syntax router [*router-address*];

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

Options ***router-address***—IP address of one or more routers.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

router-advertisement (Dynamic Profiles)

Syntax router-advertisement {...}

Hierarchy Level [edit dynamic-profiles protocols]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable router advertisement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the routing instance used to send RADIUS packets to the RADIUS server.

Options ***routing-instance-name***—Routing instance name.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Configuring the PPP Authentication Protocol*
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

routing-instance (Subscriber Secure Policy)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 15.1R3 for enhanced subscriber management on MX Series routers.

Description Specify the routing instance that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring.

Options ***routing-instance-name***—Name of the routing instance.

Default: Routing instance **default**

Required Privilege Level flow-tap—To view this statement in the configuration.
flow-tap-control—To add this statement to the configuration.

Related Documentation

- *Subscriber Secure Policy Overview*
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 579
- *Configuring Support for Subscriber Secure Policy Mirroring*

routing-instance (Diameter Peer)

Syntax routing-instance *routing-instance-name* ;

Hierarchy Level [edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

routing-instance (Diameter Transport)

Syntax routing-instance *routing-instance-name* ;

Hierarchy Level [edit diameter transport *transport-name*]

Release Information Statement introduced in Junos OS Release 11.2.

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.

Related Documentation
- Configuring the Diameter Transport on page 725

Copyright © 2017, Juniper Networks, Inc.

routing-instance-name (DHCP Local Server)

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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

routing-instance-name (DHCP Relay Agent)

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 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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

routing-instance-name (Static Subscribers)

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]

Release Information

Statement introduced in Junos OS Release 9.6.

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.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

routing-services (Enhanced Subscriber Management)

Syntax routing-services;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-name*]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Enable configuration of all routing protocols except Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on dynamically created enhanced subscriber management interfaces. IGMP and MLD are natively supported on enhanced subscriber management interfaces, and therefore do not require you to specify the **routing-services** statement in the dynamic profile.

When a dynamic profile containing the **routing-services** statement is instantiated, the router creates an enhanced subscriber management logical interface, also referred to as a pseudo logical interface, in the form demux0.*nnnnnnnnn* (for example, demux0.3221225472). Any associated subscriber routes or routes learned from a routing protocol running on the enhanced subscriber management interface use this pseudo interface as the next-hop interface.

Default If you do not include the **routing-services** statement in the dynamic profile, configuration of all routing protocols except IGMP and MLD on enhanced subscriber management 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.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Junos OS Enhanced Subscriber Management Overview on page 573

Copyright © 2017, Juniper Networks, Inc.

sdsl-bytes

Syntax sdsl-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

sdsl-overhead-adjust

Syntax sdsl-overhead-adjust *percentage*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

secret

Syntax secret *password*;

Hierarchy Level [edit access profile *profile-name* radius-server *server-address*],
[edit access radius-disconnect *client-address*],
[edit access radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.

Options ***password***—Password to use; it can include spaces if the character string is enclosed in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*
- *Configuring the RADIUS Disconnect Server for L2TP*

Copyright © 2017, Juniper Networks, Inc.

send-acct-status-on-config-change (Access Profile)

Syntax send-acct-status-on-config-change;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 13.1.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure the router's authd process to send 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.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring Per-Subscriber Session Accounting on page 105

Copyright © 2017, Juniper Networks, Inc.

send-dyn-subscription-indicator (PCRF Partition)

Syntax send-dyn-subscription-indicator;

Hierarchy Level [edit access pcrf partition *partition-name*],

Release Information Statement introduced in Junos OS Release 16.2.

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 75 on page 1359 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 75: 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

send-network-family-indicator (PCRF Partition)

Syntax send-network-family-indicator;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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 76 on page 1361.

Table 76: 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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

send-origin-state-id (Diameter Applications)

Syntax send-origin-state-id;

Hierarchy Level [edit access pcrf partition *partition-name*],
[edit access ocs partition *partition-name*],
[edit diameter peer *peer-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

send-release-on-delete (DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 10.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

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.

Related Documentation
- Extended DHCP Relay Agent Overview on page 240
- Overriding the Default DHCP Relay Configuration Settings on page 324
- Sending Release Messages When Clients Are Deleted on page 347

Copyright © 2017, Juniper Networks, Inc.

server-duid-type (DHCP Local Server)

Syntax server-duid-type *type*;

Hierarchy Level [edit system services dhcp-local-server dhcpv6]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation

- DHCPv6 Local Server Overview on page 448

Copyright © 2017, Juniper Networks, Inc.

server-group

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]

Release Information

Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

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. You can configure a maximum of five IP addresses in each named server group.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1020
- Extended DHCP Relay Agent Overview on page 240
- Configuring Named Server Groups on page 335

Copyright © 2017, Juniper Networks, Inc.

server-identifier (Address-Assignment Pools)

Syntax server-identifier *ipv4-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description 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.

Options ***ipv4-address***—IP address.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

 Copyright © 2017, Juniper Networks, Inc.

server-id-override

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]

Release Information Statement introduced in Junos OS Release 16.1.

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.

Related Documentation
- Configuring DHCP Relay Agent Source Address on page 360

Copyright © 2017, Juniper Networks, Inc.

server-response-time (DHCP Relay Agent)

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 ...]

Release Information Statement introduced in Junos OS Release 13.3.

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.

Related Documentation

- Monitoring DHCP Relay Server Responsiveness on page 384

Copyright © 2017, Juniper Networks, Inc.

service (Service Accounting)

Syntax

```
service {
  accounting-order (activation-protocol | local | radius);
  accounting {
    statistics (time | volume-time);
    update-interval minutes;
  }
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
accounting, **update-interval**, and **statistics** options added in Junos OS Release 14.2R1 for MX Series routers.

Description Define the subscriber service accounting configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Service Accounting with JSRC on page 794
- Service Accounting with JSRC on page 793
- *Configuring Service Accounting in Local Flat Files*

Copyright © 2017, Juniper Networks, Inc.

service-activation (Access Profile)

Syntax

```
service-activation {
  dynamic-profile (optional-at-login | required-at-login);
  extensible-service (optional-at-login | required-at-login);
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 16.2.

Description 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).

Options

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.

Default:

Default behavior depends on the service type:

- For **dynamic-profile** services: **required-at-login**.
- For **extensible-service** services: **optional-at-login**.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Subscriber Session Logins and Service Activation Failures Overview on page 125

Copyright © 2017, Juniper Networks, Inc.

service-context-id (OCS)

Syntax service-context-id *service-context* ;

Hierarchy Level [edit access ocs global]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring OCS Global Parameters on page 776
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756

Copyright © 2017, Juniper Networks, Inc.

- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

service-profile (DHCP Local Server)

Syntax service-profile *dynamic-profile-name*;

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* 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 ...]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Extended DHCP Local Server Overview on page 234
- Default Subscriber Service Overview on page 247
- Configuring a Default Subscriber Service on page 248

Copyright © 2017, Juniper Networks, Inc.

service-profile (DHCP Relay Agent)

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 group*group-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 ...]

Release Information
Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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 profile.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- dhcp-relay on page 1020
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- Grouping Interfaces with Common DHCP Configurations on page 287
- Default Subscriber Service Overview on page 247
- Configuring a Default Subscriber Service on page 248

Copyright © 2017, Juniper Networks, Inc.

session-options

Syntax

```
session-options {
  client-group [ group-names ];
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeoutminutes;
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description (MX Series and SRX Series devices) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Understanding Session Options for Subscriber Access on page 133
- Configuring Subscriber Session Timeout Options on page 138
- *Configuring Username Modification for Subscriber Sessions*
- Removing Inactive Dynamic Subscriber VLANs on page 139

Copyright © 2017, Juniper Networks, Inc.

shmlog (Shared Memory Log)

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]

Release Information Statement introduced in Junos OS Release 16.2.

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 ?
```

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
```


Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

-

Copyright © 2017, Juniper Networks, Inc.

sip-server-address

Syntax sip-server-address *ipv6-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family *family* dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 22. To specify multiple serv***ipv6-address***ers, add multiple **sip-server-address** statements in order of preference.

Options ***ipv6-address***—IPv6 address of a SIP outbound proxy server.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

sip-server-domain-name

Syntax sip-server-domain-name *domain-name*;

Hierarchy Level [edit access address-assignment pool *pool-name* family *family* dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the domain name of the SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 21.

Options ***domain-name***—Name of the domain.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Address-Assignment Pools Overview on page 583
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

smg-service (Enhanced Subscriber Management)

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]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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. See CLI Explorer.

Required Privilege Level
trace—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Configuring the Subscriber Management Database Trace Log Filename on page 840

Copyright © 2017, Juniper Networks, Inc.

source-address

Syntax source-address *source-address*;

Hierarchy Level [edit access radius-server *server-address*];
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for IPv6 ***source-address*** introduced in Junos OS Release 16.1.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options ***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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2017, Juniper Networks, Inc.

stacked-vlan-ranges (RADIUS Options)

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*]

Release Information Statement introduced in Junos OS Release 12.3.

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.

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

Copyright © 2017, Juniper Networks, Inc.

starts-with (DHCP Relay Agent Option)

Syntax

```
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 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 ...]

Release Information

Statement introduced in Junos OS Release 12.3.

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. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 356

Copyright © 2017, Juniper Networks, Inc.

static-subscribers

Syntax

```
static-subscribers {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      domain-name domain-name;
      interface;
      logical-system-name;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
  }
  group group-name {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include {
        domain-name domain-name;
        interface;
        logical-system-name;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
    }
    interface interface-name <exclude> <upto upto-interface-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]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Configure and associate subscribers with statically configured interfaces for dynamic service provisioning.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800

statistics (Access Profile)

Syntax statistics (time | volume-time);

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
volume-time option added in Junos OS Release 9.4.

Description Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options **time**—Collect uptime statistics only.

volume-time—Collect both volume and uptime statistics.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

statistics (Service Accounting)

Syntax statistics (time | volume-time);

Hierarchy Level [edit access profile *profile-name* service accounting]

Release Information Statement introduced in Junos OS Release 14.2R1 for MX Series routers.

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.

Related Documentation

- Configuring Service Interim Accounting on page 112
- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155

Copyright © 2017, Juniper Networks, Inc.

strict (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 450

Copyright © 2017, Juniper Networks, Inc.

strip-domain (Domain Map)

Syntax strip-domain;

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation
- Enabling Domain Name Stripping on page 174
- Configuring Domain and Realm Name Usage for Domain Maps on page 171

strip-username (Domain Map)

Syntax strip-username (*left-to-right* | *right-to-left*)

Hierarchy Level [edit access domain *domain* map]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation
- override-password on page 1244

Copyright © 2017, Juniper Networks, Inc.

subscriber-packet-idle-timeout

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*

Copyright © 2017, Juniper Networks, Inc.

subscriber-management (Subscriber Management)

Syntax

```
subscriber-management {
  enable;
  enforce-strict-scale-limit-license;
  gres-route-flush-delay;
  }
  overrides {
    no-unsolicited-ra;
    shmlog {
      disable;
      file filename <files maximum-no-files> <size maximum-file-size>;
      filtering enable;
      log-name {
        all;
        logname {
          <brief | detail | extensive | none | terse>;
          <file-logging |no-file-logging>;
        }
        }
      log-type (debug | info | notice);
    |
  }
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Router to Strictly Enforce the Subscriber Scaling License*
- *Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover*
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 302
- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Copyright © 2017, Juniper Networks, Inc.

- Configuring Junos OS Enhanced Subscriber Management on page 575

subscriber-profile

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*]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Copyright © 2017, Juniper Networks, Inc.

subscription-id-data-include (PCRF Partition)

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;
}
```

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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

Copyright © 2017, Juniper Networks, Inc.

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.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

subscription-id-type (PCRF Partition)

Syntax subscription-id-type *number*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

t1-percentage (Address-Assignment Pools)

Syntax t1-percentage *percentage*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 13.3.
DHCPv4 server support added in Junos OS Release 17.2.

Description 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.

Options *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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- DHCP Lease Timers on page 260
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

t1-renewal-time (Address-Assignment Pools)

Syntax t1-renewal-time *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family [inet | inet6] dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 17.2 on MX Series.

Description 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.

Options *seconds*—Number of seconds

Range: 30 through 4,294,967,295 seconds

Default: 50 percent of the lease duration (**preferred-lifetime**).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- DHCP Lease Timers on page 260
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

t2-percentage (Address-Assignment Pools)

Syntax t2-percentage *percentage*;

Hierarchy Level [edit access address-assignment pool *pool-name* family [inet | inet6] dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 13.3.
DHCPv4 server support added in Junos OS Release 17.2.

Description 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.

Options *percentage*—Percentage of the **preferred-lifetime** value.

Range: 0 through 100

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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- DHCP Lease Timers on page 260
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

t2-rebinding-time (Address-Assignment Pools)

Syntax t2-rebinding-time *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 17.2 on MX Series.

Description 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.

Options *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**).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- DHCP Lease Timers on page 260
- Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address on page 254
- Configuring DHCP Attributes Outside Address Assignment Pools on page 327
- DHCP Attributes for Address-Assignment Pools on page 249

Copyright © 2017, Juniper Networks, Inc.

target-logical-system (Domain Map)

Syntax

```
target-logical-system logical-system-name {
    target-routing-instance routing-instance-name;
}
```

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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 statement is explained separately. See CLI Explorer.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying a Target Logical System/Routing Instance in a Domain Map on page 170

Copyright © 2017, Juniper Networks, Inc.

target-routing-instance (Domain Map)

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*]

Release Information Statement introduced in Junos OS Release 10.4.
default option added in Junos OS Release 13.3.

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 (master) routing instance.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Understanding Domain Maps and Logical System/Routing Instance Contexts on page 163
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 170

Copyright © 2017, Juniper Networks, Inc.

terminate-code

Syntax terminate-code (aaa (deny | service-shutdown | shutdown) | dhcp | l2tp | ppp | vlan) *term-reason* radius *value*;

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 11.4.
vlan option added in Junos OS Release 16.1.

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

Copyright © 2017, Juniper Networks, Inc.

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 190
- DHCP Termination Causes and Code Values on page 192
- L2TP Termination Causes and Code Values on page 193
- PPP Termination Causes and Code Values on page 213
- VLAN Termination Causes and Code Values on page 221

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189

tftp-server

Syntax
tftp-server *ip-address*;

Hierarchy Level
[edit access address-assignment pool *pool-name* family inet dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information
Statement introduced in Junos OS Release 9.0.

Description
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.

Options
ip-address—IP address of the TFTP server.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

timeout (RADIUS)

Syntax timeout *seconds*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the amount of time that the local router or switch 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.

Options *seconds*—Amount of time to wait.

Range: 1 through 1000 seconds

Default: 3 seconds

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 113
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2017, Juniper Networks, Inc.

timeout (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 310

Copyright © 2017, Juniper Networks, Inc.

timeout (DHCP Local Server Leasequery)

Syntax timeout *seconds*;

Hierarchy Level
[edit system services dhcp-local-server allow-bulk-leasequery],
[edit system services dhcp-local-server dhcpv6 allow-bulk-leasequery],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the number of seconds that a connection on the TCP socket is 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.

Options ***seconds***—Number of seconds the connection can be idle.

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.

Related Documentation
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

timeout (DHCP Relay Agent Leasequery)

Syntax timeout *seconds*;

Hierarchy Level [edit forwarding-options dhcp-relay (leasequery | bulk-leasequery)],
[edit forwarding-options dhcp-relay dhcpv6 (leasequery | bulk-leasequery)],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* ...],
[edit routing-instances *routing-instance-name* ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the number of seconds that DHCP relay agent waits before resending a leasequery or bulk leasequery message to the DHCP servers when all servers have not responded to a previous message.

Options ***seconds***—Number of seconds.

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.

Related Documentation

- Configuring and Using DHCP Individual Leasequery on page 267
- Configuring and Using DHCP Bulk Leasequery on page 270

Copyright © 2017, Juniper Networks, Inc.

timeout-grace (RADIUS Options)

Syntax time-out grace *seconds*;

Hierarchy Level [edit access radius-options]

Release Information Statement introduced in Junos OS Release 17.2.

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.

Related Documentation

- Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable on page 121
- Global RADIUS Options for Subscriber Access on page 8
- Configuring RADIUS Server Options for Subscriber Access on page 117

Copyright © 2017, Juniper Networks, Inc.

token (DHCP Local Server)

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]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Configuring a Token for DHCP Local Server Authentication on page 371

trace (DHCP Local Server)

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*]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation

- Tracing Extended DHCP Operations on page 829
- Tracing Extended DHCP Operations for Specific Interfaces on page 834

Copyright © 2017, Juniper Networks, Inc.

trace (DHCP Relay Agent)

Syntax trace;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name*],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 10.4.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description
Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)*
- Tracing Extended DHCP Operations on page 829
- Tracing Extended DHCP Operations for Specific Interfaces on page 834

Copyright © 2017, Juniper Networks, Inc.

traceoptions (ANCP)

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]

Release Information Statement introduced in Junos OS Release 9.4.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Tracing ANCP Agent Operations for Subscriber Access on page 863

Copyright © 2017, Juniper Networks, Inc.

traceoptions (DHCP)

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]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

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.
- **interface**—Trace interface operations.
- **io**—Trace I/O operations.
- **liveness-detection**—Trace liveness detection operations.
- **packet**—Trace packet and option decoding operations.

Copyright © 2017, Juniper Networks, Inc.

- **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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Tracing Extended DHCP Operations on page 829

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Diameter Base Protocol)

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]

Release Information Statement introduced in Junos OS Release 9.6.

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.

Copyright © 2017, Juniper Networks, Inc.

- **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.

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 853

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Extensible Subscriber Services Manager)

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]**

Release Information Statement introduced in Junos OS Release 15.1.

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.

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Enhanced Subscriber Management)

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]

Release Information Statement introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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.

Copyright © 2017, Juniper Networks, Inc.

- **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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Junos OS Enhanced Subscriber Management on page 575
- Configuring the Subscriber Management Database Trace Log Filename on page 840

Copyright © 2017, Juniper Networks, Inc.

traceoptions (General Authentication Service)

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]

Release Information Statement introduced in Junos OS Release 9.0.

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

Copyright © 2017, Juniper Networks, Inc.

- **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.

Related Documentation

- Tracing General Authentication Service Processes on page 869

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Static Subscribers)

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]

Release Information

Statement introduced in Junos OS Release 9.6.

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:

Copyright © 2017, Juniper Networks, Inc.

- **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.

Related Documentation

- Tracing Static Subscriber Operations on page 881

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Subscriber Management)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.1.

Description Define tracing operations for subscriber management interface processes.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **database**—Trace database events.
- **general**—Trace general events.
- **issu**—Trace unified ISSU events.
- **server**—Trace server events.
- **session-db**—Trace session database interactions.
- **ui**—Trace user interface events.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Disable unrestricted file access.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Copyright © 2017, Juniper Networks, Inc.

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 839

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Subscriber Session Database Replication)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit system services database-replication]

Release Information Statement introduced in Junos OS Release 9.3.

Description Define tracing operations for subscriber management session database replication 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.
- **database**—Trace database events.
- **general**—Trace general flow.
- **mirror**—Trace mirroring events.
- **replication**—Trace database replication events.
- **server**—Trace server events.
- **session-db**—Trace session database interactions.
- **ui**—Trace user interface events.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

Copyright © 2017, Juniper Networks, Inc.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*****k** to specify KB, ***size*****m** to specify MB, or ***size*****g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 845

transport (Diameter Base Protocol)

Syntax

```
transport transport-name {
  address;
  logical-system logical-system-name <routing-instance routing-instance-name >;
  routing-instance routing-instance-name
}
```

Hierarchy Level
[edit diameter]

Release Information
Statement introduced in Junos OS Release 11.2.

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. See CLI Explorer.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Diameter on page 723
- Configuring the Diameter Transport on page 725

Copyright © 2017, Juniper Networks, Inc.

transport (Diameter Peer)

Syntax transport *transport-name*;

Hierarchy Level [edit diameter peer *peer-name* connect-actively]

Release Information Statement introduced in Junos OS Release 11.2.

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.

Related Documentation

- Configuring Diameter on page 723
- Configuring Diameter Peers on page 724

Copyright © 2017, Juniper Networks, Inc.

traps

Syntax (traps | no-traps);

Hierarchy Level [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*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
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.

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.

Related Documentation

- *Enabling or Disabling SNMP Notifications on Physical Interfaces*
- *Enabling or Disabling SNMP Notifications on Logical Interfaces*

Copyright © 2017, Juniper Networks, Inc.

trigger (DHCP Local Server)

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]

Release Information

Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit ... dhcpv6 ...]** hierarchy levels introduced in Junos OS Release 10.4.

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 statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 312
- radius-disconnect on page 1299

Copyright © 2017, Juniper Networks, Inc.

trigger (DHCP Relay Agent Leasequery)

Syntax trigger automatic;

Hierarchy Level [edit forwarding-options dhcp-relay dhcpv6 bulk-leasequery],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* ...],
[edit routing-instances *routing-instance-name* ...]

Release Information Statement introduced in Junos OS Release 16.1.

Description Specify the event that triggers DHCPv6 relay agent to send bulk leasequery messages to the DHCPv6 servers in the logical system/routing instance. The trigger initiates the bulk leasequery configuration parameters you specify in the DHCPv6 relay agent configuration. The optional event support is in addition to the default CLI command support.

Options **automatic**—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.

Related Documentation

- Configuring and Using DHCP Bulk Leasequery on page 270

 Copyright © 2017, Juniper Networks, Inc.

trio-flow-offload

Syntax trio-flow-offload minimum-bytes *minimum-bytes*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 12.1.

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.

Related Documentation

- *Flow Offloading*

Copyright © 2017, Juniper Networks, Inc.

trust-option-82

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]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Related Documentation
- Trusting Option 82 Information on page 346
- Overriding the Default DHCP Relay Configuration Settings on page 324

Copyright © 2017, Juniper Networks, Inc.

tunnel-profile (Domain Map)

Syntax tunnel-profile *profile-name*;

Hierarchy Level [edit access domain map *domain-map-name*]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation
- Specifying a Tunnel Profile in a Domain Map on page 175
- *Configuring a Tunnel Profile for Subscriber Access*

underlying-interface (ANCP)

Syntax underlying-interface *underlying-interface-name*;

Hierarchy Level [edit protocols ancp interfaces interface-set *interface-set-name*]

Release Information Statement introduced in Junos OS Release 12.2.

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.

Related Documentation
- Configuring the ANCP Agent on page 648
- Associating an Access Node with Subscribers for ANCP Agent Operations on page 651

Copyright © 2017, Juniper Networks, Inc.

unique-nas-port

Syntax

```
unique-nas-port {
  chassis-id chassis-id;
  chassis-id-width chassis-id-width;
}
```

Hierarchy Level [edit access radius-options]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- Enabling Unique NAS-Port Attributes (RADIUS Attribute 5) for Subscribers on page 70

 Copyright © 2017, Juniper Networks, Inc.

update-interval

Syntax update-interval *minutes*;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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 locally configured value overrides the value found in an Access-Accept message from the server.

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.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92
- Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications on page 692

Copyright © 2017, Juniper Networks, Inc.

update-interval (Service Accounting)

Syntax update-interval *minutes*;

Hierarchy Level [edit access profile *profile-name* service accounting]

Release Information Statement introduced in Junos OS Release 14.2R1 for MX Series routers.

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.

Related Documentation

- Configuring Service Interim Accounting on page 112
- Processing Cisco VSAs in RADIUS Messages for Service Provisioning on page 155
- *Configuring Service Accounting in Local Flat Files*

Copyright © 2017, Juniper Networks, Inc.

update-response-timeout (PCRF Partition)

Syntax update-response-timeout *seconds*;

Hierarchy Level [edit access pcrf partition *partition-name*]

Release Information Statement introduced in Junos OS Release 16.2.

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.

Related Documentation

- Configuring the PCRF Partition on page 773
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

Copyright © 2017, Juniper Networks, Inc.

upstream-rate (Traffic Shaping)

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]

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.

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.

Related Documentation
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 683
- Configuring the ANCP Agent on page 648
- *Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS*

Copyright © 2017, Juniper Networks, Inc.

use-interface-description

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]

Release Information
Statement introduced in Junos OS Release 9.6.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
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.

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,

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Including a Textual Description in DHCP Options on page 354
- Using DHCP Relay Agent Option 82 Information on page 349
- Configuring DHCPv6 Relay Agent Options on page 456

Copyright © 2017, Juniper Networks, Inc.

user-name-include (OCS Partition)

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*]

Release Information Statement introduced in Junos OS Release 16.2.

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

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- Configuring the OCS Partition on page 769
- 3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting on page 741
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

Copyright © 2017, Juniper Networks, Inc.

use-option-82

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)]

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.

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) (**relay-agent-remote-id** only) 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.

Related Documentation
- Using DHCP Relay Agent Option 82 Information on page 349
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 458

Copyright © 2017, Juniper Networks, Inc.

use-primary (DHCP Relay Agent)

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*]

Release Information
Statement introduced in Junos OS Release 9.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Related Documentation

- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

Copyright © 2017, Juniper Networks, Inc.

use-vlan-id

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]

Release Information
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Hierarchy level **[edit vlans *vlan-name* forwarding-options dhcp-security]** introduced in Junos OS Release 13.2X50-D10. (See *Getting Started with Enhanced Layer 2 Software* 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.

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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- *Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server*
- *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

Copyright © 2017, Juniper Networks, Inc.

user-prefix (DHCP Local Server)

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 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]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281

user-prefix (DHCP Relay Agent)

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-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]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using External AAA Authentication Services with DHCP on page 281

user-prefix (Static Subscribers)

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]

Release Information
Statement introduced in Junos OS Release 9.6.

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.

Related Documentation
- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

username-include (DHCP Local Server)

Syntax

```
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;
}
```

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 dhcpv6 authentication],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit system services dhcp-local-server group *group-name* authentication]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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:

- **client-id**
- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 281
- Creating Unique Usernames for DHCP Clients on page 282

Copyright © 2017, Juniper Networks, Inc.

username-include (DHCP Relay Agent)

Syntax

```
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;
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay 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 forwarding-options dhcp-relay group *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 ...]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
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.

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:

Copyright © 2017, Juniper Networks, Inc.

- **client-id**
- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Creating Unique Usernames for DHCP Clients on page 282
- Using External AAA Authentication Services with DHCP on page 281

Copyright © 2017, Juniper Networks, Inc.

username-include (Static Subscribers)

Syntax

```
username-include {
  domain-name domain-name;
  interface;
  logical-system-name;
  routing-instance-name;
  user-prefix user-prefix-string;
}
```

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]

Release Information

Statement introduced in Junos OS Release 9.6.

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.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 800
- Configuring the Static Subscriber Global Username on page 807
- Configuring the Static Subscriber Group Username on page 812

Copyright © 2017, Juniper Networks, Inc.

valid-lifetime (Dynamic Router Advertisement)

Syntax valid-lifetime *seconds*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols router-advertisement interface *interface-name* prefix *prefix*]

Release Information Statement introduced in Junos OS Release 10.1.

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.

Related Documentation

- *preferred-lifetime*
- Using NDRA to Provide IPv6 WAN Link Addressing Overview on page 399

Copyright © 2017, Juniper Networks, Inc.

valid-lifetime (Address-Assignment Pools)

Syntax valid-lifetime *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes],
[edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 13.3.

Description 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.

Options *seconds*—Number of seconds that the IPv6 prefix is valid.

Range: 30 through 4,294,967,295 seconds

Default: 86,400 (24 hours)

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585
- DHCP Attributes for Address-Assignment Pools on page 249
- maximum-lease-time on page 1190
- preferred-lifetime (Address-Assignment Pools) on page 1273

Copyright © 2017, Juniper Networks, Inc.

vdsl-bytes

Syntax vdsl-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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 VDSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

vdsl-overhead-adjust

Syntax vdsl-overhead-adjust *percentage*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

vdsl2-bytes

Syntax vdsl2-bytes *bytes*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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 VDSL2 access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

vdsl2-overhead-adjust

Syntax vdsl2-overhead-adjust *percentage*;

Hierarchy Level [edit protocols ancp qos-adjust]

Release Information Statement introduced in Junos OS Release 14.2.

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.

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.

Related Documentation

- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 681
- Traffic Rate Reporting and Adjustment by the ANCP Agent on page 677
- Configuring the ANCP Agent on page 648

Copyright © 2017, Juniper Networks, Inc.

vendor-specific (dhcp-relay)

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]

Release Information

Command introduced in Junos OS Release 16.2.

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.

Copyright © 2017, Juniper Networks, Inc.

violation-action (DHCP Local Server and DHCP Relay Agent)

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]

Release Information Statement introduced in Junos OS Release 14.1.

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.

Related Documentation
- Configuring a DHCP Lease-Time Threshold on page 257

Copyright © 2017, Juniper Networks, Inc.

vlan-nas-port-stacked-format

Syntax vlan-nas-port-stacked-format;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 117
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

Copyright © 2017, Juniper Networks, Inc.

vlan-ranges (RADIUS Options)

Syntax vlan-ranges (any | *low-tag–high-tag*);

Hierarchy Level [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]

Release Information Statement introduced in Junos OS Release 12.3.

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.

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 74
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 73

Copyright © 2017, Juniper Networks, Inc.

vrf-name (Duplicate Accounting)

Syntax vrf-name *vrf-name*;

Hierarchy Level [edit access profile *profile-name* accounting duplication-vrf]

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.)

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.

Related Documentation
- Understanding RADIUS Accounting Duplicate Reporting on page 99
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 92

wait-for-acct-on-ack (Access Profile)

Syntax wait-for-acct-on-ack;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 12.3.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Parameters for Subscriber Access on page 116
- Configuring Per-Subscriber Session Accounting on page 105

Copyright © 2017, Juniper Networks, Inc.

wins-server (Access)

Syntax

```
wins-server {
  ipv4-address;
}
```

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes], [edit access protocol-attributes *attribute-set-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify one or more NetBIOS name servers (NBNS) that the client uses to resolve NetBIOS names. This is equivalent to DHCP option 44.

Options ***ipv4-address***—IP address of each NetBIOS name server; add them to the configuration in order of preference.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Address-Assignment Pools on page 585

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 102

Operational Commands

- clear ancp neighbor
- clear ancp statistics
- clear ancp subscriber
- clear dhcp relay binding
- clear dhcp relay statistics
- clear dhcp server binding
- clear dhcp server statistics
- clear dhcpv6 relay binding
- clear dhcpv6 relay statistics
- clear dhcpv6 server binding
- clear dhcpv6 server statistics
- clear diameter function statistics
- clear diameter peer
- clear extensible-subscriber-services counters
- clear extensible-subscriber-services sessions
- clear ipv6 router-advertisement
- clear network-access aaa statistics
- clear network-access aaa subscriber
- clear network-access gx-plus replay
- clear network-access gx-plus statistics
- clear network-access ocs statistics
- clear network-access pcrf
- request services subscribers clear
- clear system subscriber-management statistics
- request ancp oam interface
- request ancp oam neighbor
- request ancp oam port-down
- request ancp oam port-up

Copyright © 2017, Juniper Networks, Inc.

- request dhcp relay bulk-leasequery
- request dhcp relay leasequery
- request dhcp server reconfigure
- request dhcpv6 server reconfigure
- request dhcpv6 relay bulk-leasequery
- request dhcpv6 relay leasequery
- request network-access aaa subscriber set session-id
- request services extensible-subscriber-services reload-dictionary
- request services static-subscribers login group
- request services static-subscribers logout group
- request services static-subscribers login interface
- request services static-subscribers logout interface
- request system reboot
- restart extensible-subscriber-services
- set request services subscribers
- show accounting pending-accounting-stops
- show ancp cos
- show ancp neighbor
- show ancp statistics
- show ancp subscriber
- show ancp summary
- show ancp summary neighbor
- show ancp summary subscriber
- show class-of-service interface
- show class-of-service interface-set
- show class-of-service scheduler-map
- show class-of-service traffic-control-profile
- show database-replication statistics
- show database-replication summary
- show dhcp relay binding
- show dhcp relay statistics
- show dhcp server binding
- show dhcp server statistics
- show dhcpv6 relay binding
- show dhcpv6 relay statistics
- show dhcpv6 server binding
- show dhcpv6 server statistics

Copyright © 2017, Juniper Networks, Inc.

- show diameter
- show diameter function
- show diameter function statistics
- show diameter instance
- show diameter network-element
- show diameter network-element map
- show diameter peer
- show diameter peer map
- show diameter peer statistics
- show diameter route
- show dynamic-profile session
- show extensible-subscriber-services accounting
- show extensible-subscriber-services counters
- show extensible-subscriber-services debug-information
- show extensible-subscriber-services dictionary
- show extensible-subscriber-services dictionary attributes
- show extensible-subscriber-services dictionary services
- show extensible-subscriber-services sessions
- show extensible-subscriber-services service
- show ipv6 router-advertisement
- show network-access aaa accounting
- show network-access aaa radius-servers
- show network-access aaa statistics
- show network-access aaa statistics authentication
- show network-access aaa statistics pending-accounting-stops
- show network-access aaa statistics preauthentication
- show network-access aaa subscribers
- show network-access aaa subscribers session-id
- show network-access aaa terminate-code
- show network-access address-assignment pool
- show network-access domain-map
- show network-access gx-plus
- show network-access nasreq statistics
- show network-access ocs
- show network-access pcrf
- show ppp address-pool
- show route

Copyright © 2017, Juniper Networks, Inc.

- show route extensive
- show static-subscribers sessions
- show subscribers
- show subscribers summary
- show system subscriber-management route
- show system subscriber-management statistics
- show system subscriber-management summary
- test aaa authd-lite user
- test aaa dhcp user
- test aaa ppp user

Copyright © 2017, Juniper Networks, Inc.

clear ancp neighbor

Syntax **clear ancp neighbor**
<ip-address *ip-address***>**
<system-name *mac-address***>**

Release Information Command introduced in Junos OS Release 9.4.

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

Related Documentation
- show ancp neighbor on page 1557

List of Sample Output
clear ancp neighbor on page 1475
show ancp neighbor on page 1475

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.

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

clear ancp statistics

Syntax clear ancp statistics
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information Command introduced in Junos OS Release 13.3.

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

Related Documentation
- show ancp neighbor on page 1557

List of Sample Output
clear ancp statistics on page 1477
show ancp neighbor on page 1477

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
```

Copyright © 2017, Juniper Networks, Inc.

```
            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
```


Copyright © 2017, Juniper Networks, Inc.

clear ancp subscriber

Syntax **clear ancp subscriber**
<identifier *identifier*>
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information Command introduced in Junos OS Release 11.4.

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

Related Documentation
- show ancp subscriber on page 1570

List of Sample Output
show ancp subscriber brief on page 1479
clear ancp subscriber on page 1480

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
```

Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

clear dhcp relay binding

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*>

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.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
Option **dual-stack** added in Junos OS Release 15.1.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Clearing DHCP Bindings for Subscriber Access on page 382
- show dhcp relay binding on page 1625

List of Sample Output
clear dhcp relay binding on page 1482
clear dhcp relay binding all on page 1482
clear dhcp relay binding dual-stack all on page 1482
clear dhcp relay binding interface on page 1483
clear dhcp relay binding <interfaces-vlan> on page 1483
clear dhcp relay binding <interfaces-wildcard> on page 1483

Output Fields See show dhcp relay binding 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
```


Copyright © 2017, Juniper Networks, Inc.

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.*
```

Copyright © 2017, Juniper Networks, Inc.

clear dhcp relay statistics

Syntax **clear dhcp relay statistics**
<bulk-leasequery-connections>
<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*>

Release Information Command introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options **bulk-leasequery-connections**—(Optional) Clear bulk 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

Related Documentation
- show dhcp relay statistics on page 1631

List of Sample Output clear dhcp relay statistics on page 1485

Output Fields Table 77 on page 1485 lists the output fields for the **clear dhcp relay statistics** command.

Copyright © 2017, Juniper Networks, Inc.

Table 77: clear 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.
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
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

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.

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

clear dhcp server binding

Syntax

clear dhcp server binding
<*address*>
<all>
<interface *interface-name*>
<*interfaces-vlan*>
<*interfaces-wildcard*>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

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.

Copyright © 2017, Juniper Networks, Inc.

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.

Required Privilege Level view

Related Documentation
- Clearing DHCP Bindings for Subscriber Access on page 382
- show dhcp server binding on page 1635

List of Sample Output
clear dhcp server binding <ip-address> on page 1488
clear dhcp server binding all on page 1488
clear dhcp server binding interface on page 1489
clear dhcp server binding <interfaces-vlan> on page 1489
clear dhcp server binding <interfaces-wildcard> on page 1489

Output Fields See show dhcp server binding 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:

Copyright © 2017, Juniper Networks, Inc.

```
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.*
```

Copyright © 2017, Juniper Networks, Inc.

clear dhcp server statistics

Syntax **clear dhcp server statistics**
<bulk-leasequery-connections>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 9.0.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options **bulk-leasequery-connections**—(Optional) Clear 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

List of Sample Output clear dhcp server statistics on page 1490

Output Fields See show dhcp server statistics 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
```


Copyright © 2017, Juniper Networks, Inc.

```
    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
```

Copyright © 2017, Juniper Networks, Inc.

clear dhcpv6 relay binding

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*>

Release Information

Command introduced in Junos OS Release 11.4.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
Option **dual-stack** added in Junos OS Release 15.1.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Clearing DHCP Bindings for Subscriber Access on page 382
- show dhcpv6 relay binding on page 1645

List of Sample Output
clear dhcpv6 relay binding on page 1493
clear dhcpv6 relay binding <prefix> on page 1493
clear dhcpv6 relay binding all on page 1494
clear dhcpv6 relay binding dual-stack all on page 1494
clear dhcv6p relay binding interface on page 1494
clear dhcpv6 relay binding <interfaces-vlan> on page 1494
clear dhcpv6 relay binding <interfaces-wildcard> on page 1494

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
```

Copyright © 2017, Juniper Networks, Inc.

```
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.*
```


Copyright © 2017, Juniper Networks, Inc.

clear dhcpv6 relay statistics

Syntax **clear dhcpv6 relay statistics**
<bulk-leasequery-connections>
<logical-system ***logical-system-name*****>**
<routing-instance ***routing-instance-name*****>**

Release Information Command introduced in Junos OS Release 11.4.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options **bulk-leasequery-connections**—(Optional) Clear bulk 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

List of Sample Output clear dhcpv6 relay statistics on page 1495

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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

clear dhcpv6 server binding

Syntax

clear dhcpv6 server binding
<*address*>
<all>
<interface *interface-name*>
<*interfaces-vlan*>
<*interfaces-wildcard*>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

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.

Required Privilege Level

clear

Related Documentation

- Clearing DHCP Bindings for Subscriber Access on page 382
- show dhcpv6 server binding on page 1657

Copyright © 2017, Juniper Networks, Inc.

List of Sample Output
clear dhcpv6 server binding all on page 1498
clear dhcpv6 server binding <ipv6-prefix> on page 1498
clear dhcpv6 server binding interface on page 1498
clear dhcpv6 server binding <interfaces-vlan> on page 1498
clear dhcpv6 server binding <interfaces-wildcard> on page 1498

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.*
```


Copyright © 2017, Juniper Networks, Inc.

clear dhcpv6 server statistics

Syntax

clear dhcpv6 server statistics
<bulk-leasequery-connections>
<interface *interface-name*>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections—(Optional) Clear 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

Related Documentation

- show dhcpv6 server statistics on page 1663

List of Sample Output

clear dhcpv6 server statistics on page 1499

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
```

clear diameter function statistics

Syntax **clear diameter function <*function-name*> 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.

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

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 729
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 783
- *PTSP Overview*
- show diameter on page 1666
- show diameter function on page 1672
- show diameter function statistics on page 1676

List of Sample Output clear diameter function statistics on page 1500

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear diameter function statistics

```
user@host> clear diameter function jsrc statistics
```


Copyright © 2017, Juniper Networks, Inc.

clear diameter peer

Syntax clear diameter peer *peer-name*
<connection | statistics>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation

- show diameter on page 1666
- show diameter peer on page 1687
- show diameter peer map on page 1692
- show diameter peer statistics on page 1695

List of Sample Output clear diameter peer on page 1501

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear diameter peer

```
user@host> clear diameter peer peer5 connection
```

Copyright © 2017, Juniper Networks, Inc.

clear extensible-subscriber-services counters

Syntax clear extensible-subscriber-services counters

Release Information Command introduced in Junos OS Release 15.1.

Description Clear values of all event counters to zero. Active sessions and services counters are not reset.

Required Privilege Level view

Related Documentation
- show extensible-subscriber-services counters on page 1709

Sample Output

```
#clear extensible-subscriber-services counters
```


Copyright © 2017, Juniper Networks, Inc.

clear extensible-subscriber-services sessions

Syntax clear extensible-subscriber-services sessions <*accounting-session-id*>

Release Information Command introduced in Junos OS Release 15.1.

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

Related Documentation

- show extensible-subscriber-services sessions on page 1722

Sample Output

```
#clear extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"
```

Copyright © 2017, Juniper Networks, Inc.

clear ipv6 router-advertisement

Syntax clear ipv6 router-advertisement
<interface *interface*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

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

Related Documentation
- show ipv6 router-advertisement on page 1725

List of Sample Output clear ipv6 router-advertisement on page 1504

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
```


Copyright © 2017, Juniper Networks, Inc.

clear network-access aaa statistics

Syntax

clear network-access aaa statistics
<accounting>
<address-assignment (client | pool *pool-name*)>
<authentication>
<dynamic-requests>
<radius>
<re-authentication>
<terminate-code>

Release Information

Command introduced in Junos OS Release 10.0.
Option **radius** introduced in Junos OS Release 11.4
Option **terminate-code** introduced in Junos OS Release 11.4.

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.

terminate-code—(Optional) Clear AAA termination code statistics.

Required Privilege Level

maintenance

Related Documentation

- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output

clear network-access aaa statistics accounting on page 1506
clear network-access aaa statistics address-assignment pool on page 1506
clear network-access aaa statistics radius on page 1506

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

clear network-access aaa subscriber

Syntax

clear network-access aaa subscriber
<session-id *identifier* <reconnect>>
<statistics username *username*>
<username *username* <reconnect>>

Release Information

Command introduced in Junos OS Release 9.1.
reconnect and **session-id** options added in Junos OS Release 16.1R4.

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

Related Documentation

- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output

clear network-access aaa subscriber statistics username on page 1508
clear network-access aaa subscriber username on page 1508
clear network-access aaa subscriber username on page 1508

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2017, Juniper Networks, Inc.

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 username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber session-id 18367425
```


Copyright © 2017, Juniper Networks, Inc.

clear network-access gx-plus replay

Syntax clear network-access gx-plus replay

Release Information Command introduced in Junos OS Release 11.2.

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

Related Documentation
- clear network-access gx-plus statistics on page 1510
- show network-access gx-plus on page 1770

List of Sample Output clear network-access gx-plus replay on page 1509

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
```

Copyright © 2017, Juniper Networks, Inc.

clear network-access gx-plus statistics

Syntax clear network-access gx-plus statistics

Release Information Command introduced in Junos OS Release 11.2.

Description Clear Gx-Plus statistics.

Options This command has no options.

Required Privilege Level clear

Related Documentation
- show network-access gx-plus on page 1770

List of Sample Output clear network-access gx-plus statistics on page 1510

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
```


Copyright © 2017, Juniper Networks, Inc.

clear network-access ocs statistics

Syntax clear network-access ocs statistics

Release Information Command introduced in Junos OS Release 16.2.

Description Clear Online Charging System (OCS) provisioning statistics information.

Required Privilege Level maintenance

Related Documentation

- ocs on page 1218
- show network-access ocs on page 1776
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

List of Sample Output clear network-access ocs statistics on page 1511

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
```

Copyright © 2017, Juniper Networks, Inc.

clear network-access pcrf

Syntax clear network-access pcrf (statistics | subscribers)

Release Information Command introduced in Junos OS Release 16.2.

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

Related Documentation

- pcrf on page 1259
- show network-access pcrf on page 1778
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

List of Sample Output clear network-access pcrf statistics on page 1512
clear network-access pcrf subscribers on page 1512

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
```


Copyright © 2017, Juniper Networks, Inc.

request services subscribers clear

Syntax **request services subscribers clear subscriber-profile *profile* client-id *client-id***

Release Information Command introduced in Junos OS Release 11.4.

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

List of Sample Output request services subscriber clear subscriber-profile tc_act_prof client-id on page 1513

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
```

Copyright © 2017, Juniper Networks, Inc.

clear system subscriber-management statistics

Syntax clear system subscriber-management statistics

Release Information Command introduced in Junos OS Release 15.1.

Description Clear subscriber-management statistics.

Options This command has no options.

Required Privilege Level view and system

Related Documentation
- show system subscriber-management statistics on page 1857

List of Sample Output clear subscriber-management statistics on page 1514

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
```


Copyright © 2017, Juniper Networks, Inc.

request ancp oam interface

Syntax

request ancp oam interface
(*interface-name* | interface-set *set-name*)
<count *count*>
<timeout *duration*>

Release Information

Command introduced in Junos OS Release 11.4.

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

Related Documentation

- Triggering ANCP OAM to Test the Local Loop on page 695

List of Sample Output

request ancp oam interface on page 1515

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
```

Copyright © 2017, Juniper Networks, Inc.

request ancp oam neighbor

Syntax

request ancp oam neighbor
(ip-address *ip-address* | system-name *neighbor-name*)
subscriber *identifier-string*
<count *count*>
<timeout *duration*>

Release Information

Command introduced in Junos OS Release 11.4.

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

Related Documentation

- Triggering ANCP OAM to Test the Local Loop on page 695

List of Sample Output

request ancp oam subscriber on page 1516

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”
```


Copyright © 2017, Juniper Networks, Inc.

```
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

Copyright © 2017, Juniper Networks, Inc.

request ancp oam port-down

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*

Release Information

Command introduced in Junos OS Release 16.1R4.

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

Related Documentation

- *Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages*
- *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*

List of Sample Output

request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id on page 1519

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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

request ancp oam port-up

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*

Release Information

Command introduced in Junos OS Release 16.1R4.

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

Related Documentation

- *Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages*
- *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*

List of Sample Output

request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id on page 1521

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.

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

request dhcp relay bulk-leasequery

Syntax

request dhcp relay bulk-leasequery (*client-id* | *ipv4-address* | *mac-address*)
<server-address *address* | server-group *group-name*>
<logical-system *logical-system*>
<routing-instance *routing-instance*>

Release Information

Command introduced in Junos OS Release 16.1.

Description

Initiate the bulk leasequery operation that updates binding information for subscribers by requesting that DHCP relay agent send the bulk leasequery message to the configured DHCP servers. The DHCP bulk leasequery feature must be configured in the DHCP relay agent.

Options

NOTE: You must configure at least one of *client-id*, *ipv4-address*, or *mac-address*.

address—Specify the IP address of the DHCP local server to query.

client-id—Use the DHCP client identifier option (option 61) to identify the client whose binding information is requested.

group-name—Specify the name of a group of DHCP local servers to query.

mac-address—Use the MAC address to identify the client whose binding information is requested.

logical-system—Specify an optional logical system for the DHCP server being queried. The default logical system is used by default.

routing-instance—Specify an optional routing instance for the DHCP server being queried. The default routing instance is used by default.

Required Privilege Level

view

Related Documentation

- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273

List of Sample Output

request dhcp relay bulk-leasequery on page 1523

Output Fields

When you enter this command, DHCP relay agent initiates the bulk leasequery operation.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

request dhcp relay bulk-leasequery

```
user@host> request dhcp relay bulk-leasequery 4/1224 server-address 192.168.10.10
```

Copyright © 2017, Juniper Networks, Inc.

request dhcp relay leasequery

Syntax

request dhcp relay leasequery (*client-id* | *ipv4-address* | *mac-address*)
<server-address *address* | server-group *group-name*>
<logical-system *logical-system*>
<routing-instance *routing-instance*>

Release Information

Command introduced in Junos OS Release 16.1.

Description

Initiate the leasequery operation that updates binding information for a specific subscriber by requesting that DHCP relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured in the DHCP relay agent.

Options

NOTE: You must configure at least one of *client-id*, *ipv4-address*, or *mac-address*.

address—Specify the IP address of the DHCP local server to query.

client-id—Use the DHCP client identifier option (option 61) to identify the client whose binding information is requested.

group-name—Specify the name of a group of DHCP local servers to query.

ipv4-address—Use the IPv4 address to identify the client whose binding information is requested.

logical-system—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—Specify an optional routing instance for the DHCP servers being queried. The default routing instance is used by default.

Required Privilege Level

view

Related Documentation

- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273

List of Sample Output

request dhcp relay leasequery on page 1525

Output Fields

When you enter this command, DHCP relay agent initiates the leasequery operation.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

request dhcp relay leasequery

```
user@host> request dhcp relay leasequery 192.168.25.25 server-group DHCPgroup-10
```

request dhcp server reconfigure

Syntax **request dhcp server reconfigure (all | *address* | interface *interface–name* | logical-system *logical-system–name* | routing-instance *routing-instance–name*)**

Release Information Command introduced in Junos OS Release 10.0.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

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

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

List of Sample Output request dhcp server reconfigure on page 1527

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
```

request dhcpv6 server reconfigure

Syntax **request dhcpv6 server reconfigure (all | *address* | *client-id* | interface *interface-name* | logical-system *logical-system-name* | routing-instance *routing-instance-name* | *session-id*)**

Release Information Command introduced in Junos OS Release 10.4.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

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

Related Documentation
- Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 308

List of Sample Output request dhcpv6 server reconfigure on page 1529

Output Fields When you enter this command, you are provided feedback on the status of your request.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001db8::2/16
```

Copyright © 2017, Juniper Networks, Inc.

request dhcpv6 relay bulk-leasequery

Syntax

request dhcpv6 relay bulk-leasequery ([*client-id* | *ipv6-prefix*] | relay-id *relay-id* | link-address *ipv6-link-address* | remote-id *remote-id*)
<server-address *address* | server-group *group-name*>
<logical-system *logical-system*>
<routing-instance *routing-instance*>

Release Information

Command introduced in Junos OS Release 16.1.

Description

Initiate the bulk leasequery operation that updates the DHCPv6 relay agent's binding information for a group of subscribers. The DHCPv6 bulk leasequery feature must be configured in the DHCPv6 relay agent.

By default, this bulk leasequery command is initiated with relay ID of the DHCPv6 relay agent if you do not explicitly specify one of the *client-id*, *ipv6-prefix*, *relay-id*, *ipv6-link-address*, or *remote-id* options.

Options

address—Specify the IPv6 address of the DHCP local server to query.

client-id—Update binding information for clients identified by the specified DHCPv6 Client ID option (option 1).

group-name—Specify the name of a group of DHCP local servers to query.

ipv6-link-address—Update binding information for clients associated with the specified IPv6 network segment.

ipv6-prefix—Update binding information for clients identified by the specified IPv6 prefix.

logical-system—Specify an optional logical system for the clients. The default logical system is used by default.

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.

remote-id—Update binding information for clients associated with the specified Relay Agent Remote-ID option (option 37).

routing-instance—Specify an optional routing instance for the clients. The default routing instance is used by default.

Required Privilege Level

view

Related Documentation

- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273

List of Sample Output

request dhcpv6 relay bulk-leasequery on page 1531

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

request dhcpv6 relay leasequery

Syntax

request dhcpv6 relay leasequery (*client-id* | *ipv6-prefix*)
<server-address *address* | server-group *group-name*>
<logical-system *logical-system*>
<routing-instance *routing-instance*>

Release Information

Command introduced in Junos OS Release 16.1.

Description

Initiate the leasequery operation that updates binding information for a specific subscriber by requesting that DHCPv6 relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured in the DHCPv6 relay agent.

Options

NOTE: You must configure at least one of *client-id*or *ipv6-prefix*.

address—Specify the IPv6 address of the DHCP local server to query.

client-id—Use the DHCPv6 Client ID option (option 1) to identify the client whose binding information is requested.

group-name—Specify the name of a group of DHCP local servers to query.

ipv6-prefix—Use the IPv6 prefix to identify the client whose binding information is requested.

logical-system—Specify an optional logical system for the DHCPv6 servers being queried. The default logical system is used by default.

routing-instance—Specify an optional routing instance for the DHCPv6 servers being queried. The default routing instance is used by default.

Required Privilege Level

view

Related Documentation

- Initiating DHCP Leasequery to Update the DHCP Relay Agent Lease Database on page 273

List of Sample Output

request dhcpv6 relay leasequery on page 1533

Output Fields

When you enter this command, DHCPv6 relay agent initiates the leasequery operation.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

request dhcpv6 relay leasequery

```
user@host> request dhcpv6 relay leasequery 2001:db8:3000:0:8001::5/128
```

request network-access aaa subscriber set session-id

Syntax request network-access aaa subscriber set session-id *subscriber-session-id* provisioning-state none

Release Information Command introduced in Junos OS Release 12.3.

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

Related Documentation

- Disabling PCRF Control of a Subscriber Session on page 885
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610

List of Sample Output request network-access aaa subscriber set session-id on page 1534

Output Fields When you enter this command, you are provided feedback on the status of your request. Table 78 on page 1534 lists possible error messages that might be returned if the service activation fails.

Table 78: 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.
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
```


Copyright © 2017, Juniper Networks, Inc.

```
Successful completion
```

Copyright © 2017, Juniper Networks, Inc.

request services extensible-subscriber-services reload-dictionary

Syntax **request services extensible-subscriber-services reload-dictionary**

Release Information Command introduced in Junos OS Release 15.1.

Description Reload the configured dictionary to essmd.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- dictionary on page 1043
- show extensible-subscriber-services dictionary on page 1712
- show extensible-subscriber-services dictionary attributes on page 1716
- show extensible-subscriber-services dictionary services on page 1719
- Understanding the Dictionary File on page 226

Sample Output

```
# request services extensible-subscriber-services reload-dictionary
Dictionary reloaded successfully
```


Copyright © 2017, Juniper Networks, Inc.

request services static-subscribers login group

Syntax **request services static-subscribers login group** *group-name*

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- Resetting the State of an Interface Group for Static Subscriber Login on page 820
- request services static-subscribers logout group on page 1538

List of Sample Output request services static-subscribers login group on page 1537

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
```

Copyright © 2017, Juniper Networks, Inc.

request services static-subscribers logout group

Syntax **request services static-subscribers logout group** *igroup-name*

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- Forcing a Group of Static Subscribers to Be Logged Out on page 820
- request services static-subscribers login group on page 1537

List of Sample Output request services static-subscribers logout group on page 1538

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
```


Copyright © 2017, Juniper Networks, Inc.

request services static-subscribers login interface

Syntax **request services static-subscribers login interface *interface-name***

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- Resetting the State of an Interface for Static Subscriber Login on page 819
- request services static-subscribers logout interface on page 1540

List of Sample Output request services static-subscribers login interface on page 1539

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
```

Copyright © 2017, Juniper Networks, Inc.

request services static-subscribers logout interface

Syntax **request services static-subscribers logout interface** ***interface-name***

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation

- Forcing a Static Subscriber to Be Logged Out on page 819
- request services static-subscribers login interface on page 1539

List of Sample Output request services static-subscribers logout interface on page 1540

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
```


Copyright © 2017, Juniper Networks, Inc.

request system reboot

List of Syntax

Syntax on page 1541
Syntax (EX Series Switches) on page 1541
Syntax (TX Matrix Router) on page 1541
Syntax (TX Matrix Plus Router) on page 1541
Syntax (MX Series Router) on page 1541

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)

request system reboot
<all-members>
<at *time*>
<both-routing-engines>
<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>
<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)>

Syntax (MX Series Router)

request system reboot
<all-members>
<at *time*>
<both-routing-engines>

Copyright © 2017, Juniper Networks, Inc.

<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>

Release Information Command introduced before Junos OS Release 7.4.
Option **other-routing-engine** introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **sfc** introduced for the TX Matrix Plus router in Junos OS Release 9.6.
Option **both-routing-engines** introduced in Junos OS Release 12.1.

Description Reboot the software.

NOTE: Starting with Junos OS Release 15.1F3, the statement **request system reboot** 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 statement **request system reboot** reboots only the guest operating system on the MX2010, and MX2020 with REMX2K-X8-64G.

Options **none**—Reboot the software immediately.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (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—(TX Matrix routers andTX Matrix Plus routers only) (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—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.

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.

Copyright © 2017, Juniper Networks, Inc.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at +*minutes*** option.

lcc *number*——(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis 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.

local—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.

media (compact-flash | disk)—(Optional) Boot medium for next boot.

media (external | internal)—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before stopping or rebooting the software.

other-routing-engine—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

partition—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

Copyright © 2017, Juniper Networks, Inc.

scc—(TX Matrix routers only) (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*—(TX Matrix Plus routers only) (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*—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

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 master Routing Engine, all the master 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-lcc**, **lcc *number***, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master 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 master Routing Engine.

Required Privilege Level

maintenance

Related Documentation

- *clear system reboot*

Copyright © 2017, Juniper Networks, Inc.

- *request system halt*
- *Routing Matrix with a TX Matrix Plus Router Solutions Page*
- *request vmhost reboot*

List of Sample Output

request system reboot on page 1545
request system reboot (at 2300) on page 1545
request system reboot (in 2 Hours) on page 1545
request system reboot (Immediately) on page 1545
request system reboot (at 1:20 AM) on page 1545

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
```

Copyright © 2017, Juniper Networks, Inc.

restart extensible-subscriber-services

Syntax restart extensible-subscriber-services

Release Information Command introduced in Junos OS Release 15.1.

Description Restart essmd.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- disable on page 1044

Sample Output

```
# restart extensible-subscriber-services
```


Copyright © 2017, Juniper Networks, Inc.

set request services subscribers

Syntax **request services subscribers set subscriber-profile *profile* client-id *client-id***

Release Information Command introduced in Junos OS Release 11.4.

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

List of Sample Output request services subscriber set subscriber-profile tc_act_prof client-id on page 1547

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
```

Copyright © 2017, Juniper Networks, Inc.

show accounting pending-accounting-stops

Syntax **show accounting pending-accounting-stops**
<detail | terse>
<*profile-name*>

Release Information Command introduced in Junos OS Release 13.1.

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

Related Documentation
- request network-access aaa replay pending-accounting-stops on page 1334
- show network-access aaa statistics pending-accounting-stops on page 1748

List of Sample Output
show accounting pending-accounting-stops detail on page 1550
show accounting pending-accounting-stops (Specific Profile) on page 1550
show accounting pending-accounting-stops terse on page 1551

Output Fields Table 79 on page 1548 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 79: 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

Copyright © 2017, Juniper Networks, Inc.

Table 79: show accounting pending-accounting-stops Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 79: show accounting pending-accounting-stops Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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
```

Copyright © 2017, Juniper Networks, Inc.

show ancp cos

Syntax

show ancp cos
<identifier *identifier*>
<last-update>
<pending-update>

Release Information

Command introduced in Junos OS Release 9.4.

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

Related Documentation

- show ancp neighbor on page 1557
- show ancp statistics on page 1565
- show ancp subscriber on page 1570

List of Sample Output

show ancp cos on page 1554
show ancp cos last-update on page 1555
show ancp cos pending-update on page 1556

Output Fields

Table 80 on page 1552 lists the output fields for the **show ancp cos** command. Output fields are listed in the approximate order in which they appear.

Table 80: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 80: show ancp cos Output Fields *(continued)*

Field Name	Field Description
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.
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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 80: show ancp cos Output Fields *(continued)*

Field Name	Field Description
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.
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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

show ancp neighbor

Syntax

show ancp neighbor
<brief | detail>
<ip-address ***ip-address***
<system-name ***mac-address*****>**

Release Information

Command introduced in Junos OS Release 9.4.

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

Related Documentation

- show ancp cos on page 1552
- show ancp subscriber on page 1570

List of Sample Output

show ancp neighbor on page 1560
show ancp neighbor detail on page 1561
show ancp neighbor ip-address on page 1562
show ancp neighbor system-name on page 1563

Output Fields

Table 81 on page 1557 lists the output fields for the **show ancp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 81: 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

Copyright © 2017, Juniper Networks, Inc.

Table 81: show ancp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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**

Copyright © 2017, Juniper Networks, Inc.

Table 81: 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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 81: show ancp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
* 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
```

Copyright © 2017, Juniper Networks, Inc.

```
        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
```


Copyright © 2017, Juniper Networks, Inc.

```
        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
```


Copyright © 2017, Juniper Networks, Inc.

show ancp statistics

Syntax

show ancp statistics
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information

Command introduced in Junos OS Release 13.3.

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

Related Documentation

- show ancp cos on page 1552
- show ancp neighbor on page 1557
- show ancp subscriber on page 1570

List of Sample Output

show ancp statistics on page 1567
show ancp statistics ip-address on page 1568
show ancp statistics system-name on page 1568

Output Fields

Table 82 on page 1565 lists the output fields for the **show ancp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 82: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 82: show ancp statistics Output Fields *(continued)*

Field Name	Field Description
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.
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.
Sent Generic Resp Count	Number of generic response messages sent to neighbors.

Copyright © 2017, Juniper Networks, Inc.

Table 82: show ancp statistics Output Fields *(continued)*

Field Name	Field Description
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
```

Copyright © 2017, Juniper Networks, Inc.

```
            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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show ancp subscriber

Syntax

show ancp subscriber
<brief | detail>
<identifier ***identifier*****>**
<ip-address ***ip-address*****>**
<system-name ***mac-address*****>**

Release Information

Command introduced in Junos OS Release 9.4.
neighbor option replaced with **ip-address** in Junos OS Release 16.1.
system-name option introduced in Junos OS Release 16.1.

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.

identifier ***identifier***—(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.

ip-address ***ip-address*** —(Optional) Display information about the subscribers connected to the access node specified by the IP address.

system-name ***mac-address***—(Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

Related Documentation

- clear ancp subscriber on page 1479
- show ancp cos on page 1552
- show ancp neighbor on page 1557
- show ancp statistics on page 1565

List of Sample Output

show ancp subscriber on page 1573
show ancp subscriber (After ancpd Restart) on page 1574

Copyright © 2017, Juniper Networks, Inc.

show ancp subscriber brief on page 1574
show ancp subscriber detail on page 1574
show ancp subscriber identifier identifier-string detail on page 1575

Output Fields Table 83 on page 1571 lists the output fields for the **show ancp subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 83: 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**
Type	Type of digital subscriber line employed by the access node: **ADSL1**, **ADSL2**, **ADSL2+**, **VDSL1**, **VDSL2**, **SDSL**, or **UNKNOWN**.	**brief** none
Interface	Name of the interface set or logical interface.	**brief detail** none
Rate Kbps	Actual downstream data rate for this local loop.	**brief** none
Neighbor	IP address of ANCP neighbor (access node).	**brief** none

Copyright © 2017, Juniper Networks, Inc.

Table 83: show ancp subscriber Output Fields *(continued)*

Field Name	Field Description	Level of Output
Access Loop Circuit Identifier	Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.	**detail**
Neighbor IP Address	IP address of the ANCP neighbor (access node).	**detail**
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	**detail**
DSL Type	Type of digital subscriber line employed by the access node: **ADSL1**, **ADSL2**, **ADSL2+**, **VDSL1**, **VDSL2**, **SDSL**, or **UNKNOWN**.	**detail**
DSL Line Data Link	Data link protocol employed on the access loop: **AAL5** or **Ethernet**.	**detail**
DSL Line Encapsulation	Encapsulation type on the access loop, for Ethernet only: • **0**—NA, type not conveyed • **1**—Untagged Ethernet • **2**—Single-tagged Ethernet	**detail**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 83: 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**
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  Type     Interface          Rate     Neighbor
                          State                                Kbps
**circuit 101           Idle      ADSL1    ----               32       203.0.113.13
**circuit 102           Idle      ADSL1    ----               32       203.0.113.13

  circuit 301           Showtime  ADSL1    ----               32       203.0.113.15
```

Copyright © 2017, Juniper Networks, Inc.

```
  circuit 302           Showtime  ADSL1    ----              32       203.0.113.15
```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Type     Interface         Rate     Neighbor
                         State                                Kbps
o circuit 201            Showtime  ADSL1    ----              222222
o circuit 202            Showtime  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
     DSL 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
     DSL Type                              : ADSL1
     DSL Line State                        : Idle
     DSL Line Data Link                    : Data link 2
     DSL Line Encapsulation                : N/A
```


Copyright © 2017, Juniper Networks, Inc.

```
    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 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
```

Copyright © 2017, Juniper Networks, Inc.

show ancp summary

Syntax show ancp summary

Release Information Command introduced in Junos OS Release 13.1.

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

Related Documentation

- show ancp neighbor on page 1557
- show ancp summary neighbor on page 1578
- show ancp subscriber on page 1570
- show ancp summary subscriber on page 1580

List of Sample Output show ancp summary on page 1577

Output Fields Table 84 on page 1576 lists the output fields for the **show ancp summary** command. Output fields are listed in the approximate order in which they appear.

Table 84: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 84: 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
```

Copyright © 2017, Juniper Networks, Inc.

show ancp summary neighbor

Syntax show ancp summary neighbor
<ip-address *ip-address* | system-name *mac-address*>

Release Information Command introduced in Junos OS Release 13.1.

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

Related Documentation

- show ancp summary on page 1576
- show ancp subscriber on page 1570
- show ancp summary subscriber on page 1580

List of Sample Output

show ancp summary neighbor on page 1579
show ancp summary neighbor (IP Address) on page 1579
show ancp summary neighbor (MAC Address) on page 1579

Output Fields Table 85 on page 1578 lists the output fields for the **show ancp summary** command. Output fields are listed in the approximate order in which they appear.

Table 85: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 85: show ancp summary neighbor Output Fields *(continued)*

Field Name	Field Description
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
```

Copyright © 2017, Juniper Networks, Inc.

show ancp summary subscriber

Syntax show ancp summary subscriber

Release Information Command introduced in Junos OS Release 13.1.

Description Display a summary of the counts and states for all ANCP subscribers.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show ancp summary on page 1576
- show ancp neighbor on page 1557
- show ancp summary neighbor on page 1578

List of Sample Output show ancp summary subscriber on page 1580

Output Fields Table 86 on page 1580 lists the output fields for the **show ancp summary subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 86: 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
```


Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

show class-of-service interface

Syntax show class-of-service interface
<comprehensive | detail> <*interface-name*>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Forwarding class map information added in Junos OS Release 9.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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:

Copyright © 2017, Juniper Networks, Inc.

- 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.

Required Privilege Level view

Related Documentation
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

List of Sample Output
show class-of-service interface (Physical) on page 1594
show class-of-service interface (Logical) on page 1595
show class-of-service interface (Gigabit Ethernet) on page 1595
show class-of-service interface (ANCP) on page 1595
show class-of-service interface (PPPoE Interface) on page 1595
show class-of-service interface (DHCP Interface) on page 1596
show class-of-service interface (T4000 Routers with Type 5 FPCs) on page 1596
show class-of-service interface detail on page 1596
show class-of-service interface comprehensive on page 1597
show class-of-service interface (ACX Series Routers) on page 1608
show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management) on page 1610

Output Fields Table 87 on page 1583 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 87: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
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.
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.
Congestion-notification	(QFX Series and EX4600 switches only) Congestion notification state, **enabled** or **disabled**.
Logical interface	Name of a logical interface.

Copyright © 2017, Juniper Networks, Inc.

Table 87: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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.

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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
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.
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).**
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.
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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 aborts 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.
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 aborts 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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
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.
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
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)

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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)
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)

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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**.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 87: 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.
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 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 so-0/2/3
Physical interface: so-0/2/3, Index: 135
Maximum usable queues: 8, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

  Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
```


Copyright © 2017, Juniper Networks, Inc.

```
    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 (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
```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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
  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
```


Copyright © 2017, Juniper Networks, Inc.

```
                              %           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:
```


Copyright © 2017, Juniper Networks, Inc.

```
    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
```

Copyright © 2017, Juniper Networks, Inc.

```
 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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
  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
```


Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

show class-of-service interface-set

Syntax show class-of-service interface-set
<*interface-set-name*>

Release Information Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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

List of Sample Output show class-of-service interface-set on page 1612
show class-of-service interface-set (Enhanced Subscriber Management) on page 1612

Output Fields Table 88 on page 1611 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 88: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 88: 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
```


Copyright © 2017, Juniper Networks, Inc.

```
Output traffic control profile: LEVEL_2_UID1001, Index: 4294967307
Output traffic control profile remaining: TCP_REMAIN_UID1003, Index: 4294967308
```

Copyright © 2017, Juniper Networks, Inc.

show class-of-service scheduler-map

Syntax show class-of-service scheduler-map
<*name*>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

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

Related Documentation
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

List of Sample Output
show class-of-service scheduler-map on page 1615
show class-of-service scheduler-map (QFX Series) on page 1616

Output Fields Table 89 on page 1614 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 89: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 89: show class-of-service scheduler-map Output Fields *(continued)*

Field Name	Field Description
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.
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.
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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

show class-of-service traffic-control-profile

Syntax

show class-of-service traffic-control-profile
<*profile-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series Routers.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and MPC/MIC interfaces only, display traffic shaping and scheduling profiles.

(ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.

Options

none—Display all profiles.

profile-name—(Optional) Display information about a single profile.

Required Privilege Level

view

Related Documentation

- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

List of Sample Output

show class-of-service traffic-control-profile on page 1619
show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 1620
show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 1620
show class-of-service traffic-control-profile (Enhanced Subscriber Management) on page 1620

Output Fields

Table 90 on page 1617 describes the output fields for the **show class-of-service traffic-control-profile** command. Output fields are listed in the approximate order in which they appear.

Table 90: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 90: 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.
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.

Copyright © 2017, Juniper Networks, Inc.

Table 90: show class-of-service traffic-control-profile Output Fields *(continued)*

Field Name	Field Description
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.
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

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
```

Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

show database-replication statistics

Syntax show database-replication statistics

Release Information Command introduced in Junos OS Release 9.3.

Description Display statistics regarding the replication of the subscriber management session database.

Options This command has no options.

Required Privilege Level view

Related Documentation
- show database-replication summary on page 1623

List of Sample Output show database-replication statistics on page 1621

Output Fields Table 91 on page 1621 lists the output fields for the **show database-replication statistics** command. Output fields are listed in the approximate order in which they appear.

Table 91: 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:
```

Copyright © 2017, Juniper Networks, Inc.

```
    Size (bytes)              0
    Processed                 0
Message queue:
    Queue full                0
    Max queue size            0
```


Copyright © 2017, Juniper Networks, Inc.

show database-replication summary

Syntax show database-replication summary

Release Information Command introduced in Junos OS Release 9.3.

Description Display summary information regarding database replication for the subscriber management session database.

Options This command has no options.

Required Privilege Level view

Related Documentation
- show database-replication statistics on page 1621

List of Sample Output show database-replication summary on page 1624

Output Fields Table 92 on page 1623 lists the output fields for the **show database-replication summary** command. Output fields are listed in the approximate order in which they appear.

Table 92: 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: • **Master** • **Standby**
Connection	State of the connection: • **Up** • **Down**
Database	State of the subscriber management database: • **Available** • **Unavailable** • **Synchronized**

Copyright © 2017, Juniper Networks, Inc.

Table 92: 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
```


Copyright © 2017, Juniper Networks, Inc.

show dhcp relay binding

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>

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.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

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.

Copyright © 2017, Juniper Networks, Inc.

summary—(Optional) Display a summary of DHCP client information.

Required Privilege Level view

Related Documentation
- Clearing DHCP Bindings for Subscriber Access on page 382
- clear dhcp relay binding on page 1481

List of Sample Output
show dhcp relay binding on page 1628
show dhcp relay binding detail on page 1628
show dhcp relay binding interface on page 1628
show dhcp relay binding interface vlan-id on page 1629
show dhcp relay binding interface svlan-id on page 1629
show dhcp relay binding ip-address on page 1629
show dhcp relay binding mac-address on page 1629
show dhcp relay binding session-id on page 1629
show dhcp relay binding <interfaces-vlan> on page 1629
show dhcp relay binding <interfaces-wildcard> on page 1629
show dhcp relay binding summary on page 1630

Output Fields Table 93 on page 1626 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 93: 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

Copyright © 2017, Juniper Networks, Inc.

Table 93: show dhcp relay binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
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

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

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)
```


Copyright © 2017, Juniper Networks, Inc.

show dhcp relay statistics

Syntax

show dhcp relay statistics
<**bulk-leasequery-connections**>
<**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*>

Release Information

Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections—(Optional) Display information about bulk leasequery operations.

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

Related Documentation

- clear dhcp relay statistics on page 1484

List of Sample Output

show dhcp relay statistics on page 1633
show dhcp relay statistics bulk-leasequery-connections on page 1634

Output Fields

Table 94 on page 1632 lists the output fields for the **show dhcp relay statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 94: 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.
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

Copyright © 2017, Juniper Networks, Inc.

Table 94: show dhcp relay statistics Output Fields *(continued)*

Field Name	Field Description
External Server Response	State of the external DHCP server responsiveness.
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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

show dhcp server binding

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*>

Release Information

Command introduced in Junos OS Release 9.0.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation

- Clearing DHCP Bindings for Subscriber Access on page 382
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- clear dhcp server binding on page 1487

List of Sample Output

show dhcp server binding on page 1638
show dhcp server binding detail on page 1639
show dhcp server binding detail (ACI Interface Set Configured) on page 1639
show dhcp server binding interface <vlan-id> on page 1640
show dhcp server binding interface <svlan-id> on page 1640
show dhcp server binding <ip-address> on page 1640
show dhcp server binding <session-id> on page 1640
show dhcp server binding summary on page 1640
show dhcp server binding <interfaces-vlan> on page 1640
show dhcp server binding <interfaces-wildcard> on page 1640

Output Fields Table 95 on page 1636 lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 95: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 95: show dhcp server binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 95: 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. 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.	**detail**
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**

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

show dhcp server statistics

Syntax

show dhcp server statistics
<**bulk-leasequery-connections**>
<**logical-system** *logical-system-name*>
<**routing-instance** *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections—(Optional) Display bulk leasequery statistics.

bulk-leasequery-connections—(Optional) Display information about 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

Related Documentation

- clear dhcp server statistics on page 1490

List of Sample Output

show dhcp server statistics on page 1643
show dhcp server statistics on page 1644

Output Fields

Table 96 on page 1642 lists the output fields for the **show dhcp server statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 96: 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
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.

Copyright © 2017, Juniper Networks, Inc.

Table 96: 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

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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

show dhcpv6 relay binding

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>

Release Information

Command introduced in Junos OS Release 11.4.
interfaces-vlan and *interfaces-wildcard* options introduced in Junos OS Release 12.1.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Clearing DHCP Bindings for Subscriber Access on page 382
- clear dhcpv6 relay binding on page 1492

List of Sample Output
show dhcpv6 relay binding on page 1648
show dhcpv6 relay binding (Address) on page 1648
show dhcpv6 relay binding detail (Client ID) on page 1649
show dhcpv6 relay binding detail on page 1649
show dhcpv6 relay binding detail (Dual-Stack) on page 1649
show dhcpv6 relay binding detail (Multi-Relay Topology) on page 1650
show dhcpv6 relay binding (Session ID) on page 1650
show dhcpv6 relay binding (Subscriber with Multiple Addresses) on page 1650
show dhcpv6 relay binding detail (Subscriber with Multiple Addresses) on page 1651
show dhcpv6 relay binding (Interfaces VLAN) on page 1651
show dhcpv6 relay binding (Interfaces Wildcard) on page 1652
show dhcpv6 relay binding (Interfaces Wildcard) on page 1652
show dhcpv6 relay binding summary on page 1652

Output Fields Table 97 on page 1646 lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 97: 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 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

Copyright © 2017, Juniper Networks, Inc.

Table 97: 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**
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

Copyright © 2017, Juniper Networks, Inc.

Table 97: show dhcpv6 relay binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
     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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

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)
```


Copyright © 2017, Juniper Networks, Inc.

show dhcpv6 relay statistics

Syntax

show dhcpv6 relay statistics
<bulk-leasequery-connections>
<logical-system ***logical-system-name*****>**
<routing-instance ***routing-instance-name*****>**

Release Information

Command introduced in Junos OS Release 11.4.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

bulk-leasequery-connections—(Optional) Display bulk 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

Related Documentation

- clear dhcpv6 relay statistics on page 1495

List of Sample Output

show dhcpv6 relay statistics on page 1655
show dhcpv6 relay statistics bulk-leasequery-connections on page 1656

Output Fields

Table 98 on page 1654 lists the output fields for the **show dhcpv6 relay statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 98: 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. • **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.
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_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 • **DHCP6_LEASEQUERY**—Number of DHCP leasequery messages transmitted
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.

Copyright © 2017, Juniper Networks, Inc.

Table 98: show dhcpv6 relay statistics Output Fields *(continued)*

Field Name	Field Description
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                       1
    Lease Time Violation        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
    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_FORW           0
    DHCPV6_LEASEQUERY           0

Packets forwarded:
    Total                       4
    FWD REQUEST                 2
```

Copyright © 2017, Juniper Networks, Inc.

```
    FWD REPLY                   2

External Server Response:
    State                      Responding
```

show dhcpv6 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
```


Copyright © 2017, Juniper Networks, Inc.

show dhcpv6 server binding

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*>

Release Information

Command introduced in Junos OS Release 9.6.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

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

Related Documentation

- Clearing DHCP Bindings for Subscriber Access on page 382
- clear dhcpv6 server binding on page 1497

Copyright © 2017, Juniper Networks, Inc.

List of Sample Output

show dhcpv6 server binding on page 1659
show dhcpv6 server binding detail on page 1660
show dhcpv6 server binding interface on page 1660
show dhcpv6 server binding interface detail on page 1660
show dhcpv6 server binding (IPv6 Prefix) on page 1661
show dhcpv6 server binding (Session ID) on page 1661
show dhcpv6 server binding (Interfaces VLAN) on page 1661
show dhcpv6 server binding (Interfaces Wildcard) on page 1661
show dhcpv6 server binding (Interfaces Wildcard) on page 1662
show dhcpv6 server binding summary on page 1662

Output Fields Table 99 on page 1658 lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

Table 99: 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**
Client IPv6 Prefix	Client's IPv6 prefix.	**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**

Copyright © 2017, Juniper Networks, Inc.

Table 99: show dhcpv6 server binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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

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
```

Copyright © 2017, Juniper Networks, Inc.

show dhcpv6 server binding detail

```
user@host> show dhcpv6 server binding detail
Session Id:  6
     Client IPv6 Prefix:                  2001:db8:1111:2222::/64
     Client DUID:                        LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

     State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                       2009-07-21 10:41:15 PDT
     Lease Expires in:                    86308 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
     Lease time violated:                 yes
     Incoming Client Interface:           ge-1/0/0.0
     Server Ip Address:                   0.0.0.0
     Server Interface:                    none
     Client Id Length:                    14
     Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0001

Session Id:  7
     Client IPv6 Address:                 2001:db8:2002::1/128
     Client IPv6 Prefix:                  2001:db8:1111:2222::/64
     Client DUID:                        LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

     State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                       2009-07-21 10:41:15 PDT
     Lease Expires in:                    86308 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
     Client Pool Name:                    bos-v6-pool
     Client Prefix Pool Name:             bos-v6-prefix-pool
     Client Id Length:                    14
     Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002
```

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
```


Copyright © 2017, Juniper Networks, Inc.

```
     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
```

Copyright © 2017, Juniper Networks, Inc.

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)
```


Copyright © 2017, Juniper Networks, Inc.

show dhcpv6 server statistics

Syntax

show dhcpv6 server statistics
<bulk-leasequery-connections>
<logical-system ***logical-system-name*****>**
<routing-instance ***routing-instance-name*****>**

Release Information

Command introduced in Junos OS Release 9.6.
bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections—(Optional) Display 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

Related Documentation

- clear dhcpv6 server statistics on page 1499

List of Sample Output

show dhcpv6 server statistics on page 1665
show dhcpv6 server statistics bulk-leasequery-connections on page 1665

Output Fields

Table 100 on page 1664 lists the output fields for the **show dhcpv6 server statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 100: 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 • **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
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.

Copyright © 2017, Juniper Networks, Inc.

Table 100: 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_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                     1
    Lease Time Violation      1

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
    DHCPV6_RELAY_REPL         0
    DHCPV^_LEASEQUERY         0

Messages sent:
    DHCPV6_ADVERTISE          9
    DHCPV6_REPLY              5
    DHCPV6_RECONFIGURE        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
```

Copyright © 2017, Juniper Networks, Inc.

show diameter

Syntax **show diameter**
<brief | detail | summary>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation

- clear diameter function statistics on page 1500
- clear diameter peer on page 1501
- show diameter function on page 1672
- show diameter instance on page 1679
- show diameter network-element on page 1681
- show diameter peer on page 1687
- show diameter route on page 1699

List of Sample Output

show diameter brief on page 1669
show diameter detail on page 1670
show diameter summary on page 1670

Output Fields Table 101 on page 1666 lists the output fields for the **show diameter** command. Output fields are listed in the approximate order in which they appear.

Table 101: 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

Copyright © 2017, Juniper Networks, Inc.

Table 101: show diameter Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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.	**brief detail**

Copyright © 2017, Juniper Networks, Inc.

Table 101: 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).	**brief detail**
Upstream Transaction Utilization	Percent of upstream traffic used for this function.	**brief detail**
Downstream Transaction Utilization	Percent of downstream traffic used for this function.	**brief detail**
Net Queue Buffer Utilization	Percent of network transmission buffer used for this function.	**brief detail**
Func Queue Buffer Utilization	Percent of function transmission buffer used for this function.	**brief detail**
Routed Dests	Number of destinations that have this function associated with their routes.	**brief detail**
Name	Name of the Diameter instance.	**brief detail**
Origin-Realm	Value of Origin-Realm attribute-value pair (AVP).	**brief detail**
Origin Host	Value of Origin-Host AVP.	**brief detail**
NE-Total	Number of configured network elements.	**brief detail**
NE-Connected	Number of network elements with active Diameter connections.	**brief detail**
Name	Name of the Diameter network element.	**brief detail**
Instance	Name of the Diameter instance in which the network element is configured.	**brief detail**
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.	**brief detail**
Primary Peer	Primary peer for the network element, based on the configured peer priority.	**brief detail**
Secondary Peer	Secondary peer for the network element, based on the configured peer priority.	**brief detail**

Copyright © 2017, Juniper Networks, Inc.

Table 101: show diameter Output Fields *(continued)*

Field Name	Field Description	Level of Output
Peer	Name of the peer.	**brief detail**
Instance	Name of the Diameter instance in which the peer is configured.	**brief detail**
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.	**brief detail**
NE-Count	Number of network elements associated with the peer.	**brief detail**
Activated Count	Activation status of the peer: • **1**—Peer is activated. • **0**—Peer is not activated.	**brief detail**
Primary Count	Status of the peer: primary (**1**) or secondary (**0**).	**brief detail**
Secondary Count	Status of the peer: secondary (**0**) or primary (**1**).	**brief detail**
Route	Name of the Diameter route.	**detail**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show diameter function

Syntax

show diameter function
<brief | detail | summary>
<function-name>

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.

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

Related Documentation

- clear diameter function statistics on page 1500
- show diameter on page 1666
- show diameter function statistics on page 1676

List of Sample Output

show diameter function on page 1674
show diameter function brief on page 1674
show diameter function detail (JSRC) on page 1675
show diameter function detail (Gx-Plus) on page 1675

Output Fields

Table 102 on page 1672 lists the output fields for the **show diameter function** command. Output fields are listed in the approximate order in which they appear.

Table 102: 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

Copyright © 2017, Juniper Networks, Inc.

Table 102: show diameter function Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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**

Copyright © 2017, Juniper Networks, Inc.

Table 102: 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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

show diameter function statistics

Syntax

show diameter function statistics
<brief | detail | summary>
<function-name>

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.

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

Related Documentation

- clear diameter function statistics on page 1500
- show diameter on page 1666
- show diameter function on page 1672

List of Sample Output

show diameter function statistics on page 1678
show diameter function statistics brief on page 1678
show diameter function statistics detail on page 1678

Output Fields

Table 103 on page 1676 lists the output fields for the **show diameter function statistics** command. Output fields are listed in the approximate order in which they appear.

Table 103: 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

Copyright © 2017, Juniper Networks, Inc.

Table 103: show diameter function statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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**
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**

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

show diameter instance

Syntax **show diameter instance**
<brief | detail | summary>
<instance-name>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- show diameter on page 1666

List of Sample Output
show diameter instance on page 1680
show diameter instance detail on page 1680

Output Fields Table 104 on page 1679 lists the output fields for the **show diameter instance** command. Output fields are listed in the approximate order in which they appear.

Table 104: 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**

Copyright © 2017, Juniper Networks, Inc.

Table 104: show diameter instance Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```


Copyright © 2017, Juniper Networks, Inc.

show diameter network-element

Syntax

show diameter network-element
<brief | detail | summary>
<element-name>

Release Information

Command introduced in Junos OS Release 9.6.

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

Related Documentation

- show diameter on page 1666
- show diameter function on page 1672
- show diameter network-element map on page 1684
- show diameter peer on page 1687
- show diameter route on page 1699

List of Sample Output

show diameter network-element on page 1682
show diameter network-element detail on page 1683

Output Fields

Table 105 on page 1681 lists the output fields for the **show diameter network-element** command. Output fields are listed in the approximate order in which they appear.

Table 105: 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**

Copyright © 2017, Juniper Networks, Inc.

Table 105: show diameter network-element Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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
```

Copyright © 2017, Juniper Networks, Inc.

show diameter network-element map

Syntax **show diameter network-element map**
<brief | detail | summary>
<element-name>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- show diameter on page 1666
- show diameter network-element on page 1681

List of Sample Output
show diameter network-element map on page 1685
show diameter network-element map detail on page 1685

Output Fields Table 106 on page 1684 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 106: 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

Copyright © 2017, Juniper Networks, Inc.

Table 106: show diameter network-element map Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

show diameter peer

Syntax

show diameter peer
<brief | detail | summary>
<peer-name>

Release Information

Command introduced in Junos OS Release 9.6.

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

Related Documentation

- clear diameter peer on page 1501
- show diameter on page 1666
- show diameter peer map on page 1692
- show diameter peer statistics on page 1695

List of Sample Output

show diameter peer on page 1689
show diameter peer detail on page 1690

Output Fields

Table 107 on page 1687 lists the output fields for the **show diameter peer** command. Output fields are listed in the approximate order in which they appear.

Table 107: 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**

Copyright © 2017, Juniper Networks, Inc.

Table 107: show diameter peer 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. • 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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 107: show diameter peer Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local transport	Number of transports configured.	**detail**
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**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
Suspended timeout base  :     30000 ms
Closed timeout          :     30000 ms
Connection timeout      :      6000 ms
Waiting origin state id : no
```

Copyright © 2017, Juniper Networks, Inc.

show diameter peer map

Syntax **show diameter peer map**
<brief | detail | summary>
<peer-name>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- clear diameter peer on page 1501
- show diameter on page 1666
- show diameter peer on page 1687
- show diameter peer statistics on page 1695

List of Sample Output
show diameter peer map on page 1693
show diameter peer map detail on page 1693

Output Fields Table 108 on page 1692 lists the output fields for the **show diameter peer map** command. Output fields are listed in the approximate order in which they appear.

Table 108: 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
NE	Name of the Diameter network element.	All levels
Priority	Priority configured for the peer. A lower number indicates a higher priority.	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 108: show diameter peer map Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

show diameter peer statistics

Syntax

show diameter peer statistics
<brief | detail | summary>
<peer-name>

Release Information

Command introduced in Junos OS Release 9.6.

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

Related Documentation

- clear diameter peer on page 1501
- show diameter on page 1666
- show diameter peer on page 1687
- show diameter peer map on page 1692

List of Sample Output

show diameter peer statistics on page 1696
show diameter peer statistics detail on page 1696

Output Fields

Table 109 on page 1695 lists the output fields for the **show diameter peer statistics** command. Output fields are listed in the approximate order in which they appear.

Table 109: 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**

Copyright © 2017, Juniper Networks, Inc.

Table 109: show diameter peer statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

show diameter route

Syntax
show diameter route
<brief | detail | summary>
<route-name>

Release Information
Command introduced in Junos OS Release 9.6.

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

Related Documentation
- show diameter on page 1666
- show diameter network-element on page 1681

List of Sample Output
show diameter route on page 1700
show diameter route detail on page 1700

Output Fields
Table 110 on page 1699 lists the output fields for the **show diameter route** command. Output fields are listed in the approximate order in which they appear.

Table 110: 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

Copyright © 2017, Juniper Networks, Inc.

Table 110: 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
```


Copyright © 2017, Juniper Networks, Inc.

show dynamic-profile session

Syntax

show dynamic-profile session
<client-id *client-id*>
<profile-name *profile-name*>
<service-id *service-id*>

Release Information

Command introduced in Junos OS Release 13.3.

Description

Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.

NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as NONE in the command output.
- The variable is also displayed as NONE when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the **show configuration dynamic-profiles** command.

Options

client-id ***client-id***—Display dynamic profile information for subscribers associated with the specified client.

profile-name ***profile-name***—(Optional) Display dynamic profile information for the specified subscriber or service profile.

service-id ***service-id***—Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level

view

List of Sample Output

show dynamic-profile session client-id (Client ID) on page 1702
show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile) on page 1704

Copyright © 2017, Juniper Networks, Inc.

show dynamic-profile session service-id (Service Session) on page 1704

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 {
```


Copyright © 2017, Juniper Networks, Inc.

```
    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;
            }
        }
    }
```

Copyright © 2017, Juniper Networks, Inc.

```
}
bsimmons
}
```

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile)

```
user@host>show dynamic-profile session client-id 20 profile-name cos-service
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session service-id (Service Session)

```
user@host>show dynamic-profile session service-id 21
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
```


Copyright © 2017, Juniper Networks, Inc.

```
                term rest {
                    then accept;
                }
            }
        }
    }
    policer policer1_UID1025 {
        if-exceeding {
            bandwidth-limit 1m;
            burst-size-limit 15k;
        }
        then discard;
    }
  }
}
```

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services accounting

Syntax show extensible-subscriber-services accounting vsa *<vsa>*

Release Information Command introduced in Junos OS Release 15.1.

Description Display the accounting statistics collected in the current log file. If the **vsa** option is not provided, the entire content of the current statistics file is displayed. If the **vsa** option and value are specified, contents of each line that matches the vendor-specific attribute are displayed.

Options *vsa*—(Optional) Vendor-specific attribute.

Range: 1 through 255

Required Privilege Level view

Output Fields Table 111 on page 1706 lists the output fields for the **show extensible-subscriber-services accounting** command. Output fields are listed in the approximate order in which they appear.

Table 111: show extensible-subscriber-services accounting Output Fields

Field Name	Field Description
Timestamp	System time when the statistics file is generated in UTC.
Accounting-Status-Type	Type of status, whether it is service start, stop, or update message
NAS-Port-ID	NAS port ID
Line-ID	Line ID
Vlan-ID	VLAN ID
Logical-Interface	Physical and logical interface information
Total-Input-Bytes	Number of bytes received on the interface
Total-Input-Packets	Number of packets received on the interface
IPv6-Input-Bytes	Number of IPv6 bytes received on the interface
IPv6-Input-Packets	Number of IPv6 packets received on the interface
Total-Output-Bytes	Number of bytes transmitted from the interface
Total-Output-Packets	Number of packets transmitted from the interface

Copyright © 2017, Juniper Networks, Inc.

Table 111: show extensible-subscriber-services accounting Output Fields *(continued)*

Field Name	Field Description
IPv6-Output-Bytes	Number of IPv6 bytes transmitted from the interface
IPv6-Output-Packets	Number of IPv6 packets transmitted from the interface
Egress-Queued-Packets	Number of packets sent for this class
Egress-Queued-Bytes	Number of bytes, sent for this class
Egress-TX-Packets	Number of packets transmitted by this queue
Egress-TX-Bytes	Number of bytes actually transmitted by this queue
Egress-Tail-Drop-Packets	Number of packets dropped because of tail drop
Egress-Red-Drop-Packets	Number of packets dropped because of random early detection
Egress-Red-Drop-Bytes	Number of bytes dropped because of random early detection

Sample Output

```
#show extensible-subscriber-services accounting  vsa
<Ipdr>
  <General-Interface-Parameters>
        <Timestamp>2013-02-19T01:23:51</Timestamp>
        <Accounting-Status-Type>start</Accounting-Status-Type>

<Description>#ge-1/0/0.demux0.1073770301:635-7#ngcoco#remote634#circuit634#</Description>


<NAS-Port-ID>mynah#ge-1/0/0.demux0.1073770301:635-7#circuit634#remote634</NAS-Port-ID>

        <Line-ID>remote634</Line-ID>
        <Vlan-ID>635</Vlan-ID>
        <Logical-Interface>ge-1/0/0.2395</Logical-Interface>
  </General-Interface-Parameters>
  <Overall-Packets>
        <Total-Input-Bytes>0</Total-Input-Bytes>
        <Total-Input-Packets>0</Total-Input-Packets>
        <IPv6-Input-Bytes>0</IPv6-Input-Bytes>
        <IPv6-Input-Packets>0</IPv6-Input-Packets>
        <Total-Output-Bytes>0</Total-Output-Bytes>
        <Total-Output-Packets>0</Total-Output-Packets>
        <IPv6-Output-Bytes>0</IPv6-Output-Bytes>
        <IPv6-Output-Packets>0</IPv6-Output-Packets>
  </Overall-Packets>
  <Egress-Packets>
        <Egress-Queue-Index>0</Egress-Queue-Index>
        <Egress-Queued-Packets>0</Egress-Queued-Packets>
        <Egress-Queued-Bytes>0</Egress-Queued-Bytes>
        <Egress-TX-Packets>0</Egress-TX-Packets>
        <Egress-TX-Bytes>0</Egress-TX-Bytes>
```

Copyright © 2017, Juniper Networks, Inc.

```
      <Egress-Tail-Drop-Packets>0</Egress-Tail-Drop-Packets>
      <Egress-Red-Drop-Packets>0</Egress-Red-Drop-Packets>
      <Egress-Red-Drop-Bytes>0</Egress-Red-Drop-Bytes>
</Egress-Packets>
<Egress-Packets>
      <Egress-Queue-Index>1</Egress-Queue-Index>
      <Egress-Queued-Packets>0</Egress-Queued-Packets>
      <Egress-Queued-Bytes>0</Egress-Queued-Bytes>
      <Egress-TX-Packets>0</Egress-TX-Packets>
      <Egress-TX-Bytes>0</Egress-TX-Bytes>
      <Egress-Tail-Drop-Packets>0</Egress-Tail-Drop-Packets>
      <Egress-Red-Drop-Packets>0</Egress-Red-Drop-Packets>
      <Egress-Red-Drop-Bytes>0</Egress-Red-Drop-Bytes>
</Egress-Packets>
<Egress-Packets>
      <Egress-Queue-Index>2</Egress-Queue-Index>
      <Egress-Queued-Packets>0</Egress-Queued-Packets>
      <Egress-Queued-Bytes>0</Egress-Queued-Bytes>
      <Egress-TX-Packets>0</Egress-TX-Packets>
      <Egress-TX-Bytes>0</Egress-TX-Bytes>
      <Egress-Tail-Drop-Packets>0</Egress-Tail-Drop-Packets>
      <Egress-Red-Drop-Packets>0</Egress-Red-Drop-Packets>
      <Egress-Red-Drop-Bytes>0</Egress-Red-Drop-Bytes>
</Egress-Packets>
```


Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services counters

Syntax show extensible-subscriber-services counters

Release Information Command introduced in Junos OS Release 15.1.

Description Display number of times various events were processed by Extensible Subscriber Services Manager.

Required Privilege Level view

Related Documentation
- clear extensible-subscriber-services counters on page 1502

Output Fields Table 112 on page 1709 describes the output fields for the **show extensible-subscriber-services counters** command. Output fields are listed in the approximate order in which they appear.

Table 112: show extensible-subscriber-services counters Output Fields

Field Name	Field Description
Total Sessions	Number of extensible-subscriber-service control sessions that are currently in the system
Total Services	Number of extensible-subscriber-services that are currently in the system
Active Services	Number of extensible-subscriber-services that are currently in active state
Op-script execution	Total number of operational scripts executed: • Successful—Operational scripts that are executed without error • Unsuccessful—Operational scripts that encountered errors
Commit execution	Total number of commit scripts executed: • Successful—Commit scripts that are executed without error • Unsuccessful—Commit scripts that encountered errors
Application execution	Number of applications executed: • Successful—Applications that are executed without error • Unsuccessful—Applications that encountered errors
Service requests	Number of service requests received by the daemon. It has counters for total number of requests received, acknowledged (Acked), and negatively acknowledged (Nacked).

Copyright © 2017, Juniper Networks, Inc.

Sample Output

```
# show extensible-subscriber-services counters

Total Sessions:         7
Total Services:         14
Active Services:        14

Op-script execution:
    Successful          30
    Unsuccessful        0
Commit execution:
    Successful          4
    Unsuccessful        0
Application execution:
    Successful          0
    Unsuccessful        0
Service requests:
    Received            30
    Acked               30
    Nacked              0
```


Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services debug-information

Syntax show extensible-subscriber-services debug-information

Release Information Command introduced in Junos OS Release 15.1.

Description Write service session information to a file.

Required Privilege Level view

Output Fields Table 113 on page 1711 lists the output fields for the **show extensible-subscriber-services debug-information** command.

Table 113: show extensible-subscriber-services debug-information Output Fields

Field Name	Field Description
File Name	Name of the file that contains the service session information

Sample Output

```
#show extensible-subscriber-services debug-information
File Name: /var/tmp/dump_session_record
```

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services dictionary

Syntax show extensible-subscriber-services dictionary

Release Information Command introduced in Junos OS Release 15.1.

Description Display contents of the dictionary including attributes and services.

Required Privilege Level view

Related Documentation

- dictionary on page 1043
- request services extensible-subscriber-services reload-dictionary on page 1536
- show extensible-subscriber-services dictionary attributes on page 1716
- show extensible-subscriber-services dictionary services on page 1719
- Understanding the Dictionary File on page 226

Output Fields Table 114 on page 1712 lists the output fields for the **show extensible-subscriber-services dictionary** command. Output fields are listed in the approximate order in which they appear.

Table 114: show extensible-subscriber-services dictionary Output Fields

Field Name	Field Description
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**
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.

Copyright © 2017, Juniper Networks, Inc.

Table 114: show extensible-subscriber-services dictionary Output Fields *(continued)*

Field Name	Field Description
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	Traffic mirroring action. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26–59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26–58, 26–59, 26–60, or 26–61) must always include all four VSAs. If the CoA action is to stop mirroring (VSA 26–58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary

RADIUS DICTIONARY

  Attribute Name             Vendor ID   Attribute Code    Attribute Type    Has Tag   Config Type
  Class                      0           25                string            FALSE     INVALID

  Acct-Session-Id            0           44                string            FALSE     INVALID

  ERX-Service-Activate       4874        65                string            TRUE      SERVICE_ACTIVATE

  Service-Type               4874        173               integer           TRUE      INVALID

  ADSL-Agent-Remote-Id       3561        2                 string            FALSE     INVALID

  ERX-Med-Port-Number        4874        61                OctetString       FALSE     INVALID

  NAS-Port-Id                0           87                text              FALSE     INVALID

  ERX-Med-Dev-Handle         4874        59                OctetString       FALSE     INVALID

  ERX-Service-Deactivate     4874        66                string            TRUE      SERVICE_DEACTIVATE

  NAS-IP-Address             0           4                 address           FALSE     INVALID

  NAS-Identifier             0           32                string            FALSE     INVALID
```

Copyright © 2017, Juniper Networks, Inc.

```
ERX-Service-Acct-Interval   4874        140              integer          TRUE      INVALID

ADSL-Agent-Circuit-Id       3561        1                string           FALSE     INVALID

ERX-LI-Action               4874        58               OctetString      FALSE     SERVICE_ACTIVATE

ERX-Med-Ip-Address          4874        60               OctetString      FALSE     INVALID


Services List

Service Name: ngcoco
  Service Attribute Name      : ERX-Service-Activate

  PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_ngcoco_add

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

Service Name: dhcprelay
  Service Attribute Name      : ERX-Service-Activate

  PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_dhcprelay_add

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

Service Name: default
  Service Attribute Name      : ERX-LI-Action

  PROVISION ACTION
```


Copyright © 2017, Juniper Networks, Inc.

```
    Action Type               : APPLICATION
    Action Version            : 1
    Action Name               : LI

    PARAMETERS
      ERX-Med-Dev-Handle
      ERX-Med-Ip-Address
      ERX-Med-Port-Number

Service Name: dhcprelay
  Service Attribute Name      : ERX-Service-Deactivate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Deactivate

Service Name: ngcoco
  Service Attribute Name      : ERX-Service-Deactivate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Deactivate
```

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services dictionary attributes

Syntax show extensible-subscriber-services dictionary attributes

Release Information Command introduced in Junos OS Release 15.1.

Description Display contents of the dictionary file. It shows only the parameter attributes in the dictionary.

Required Privilege Level view

Related Documentation

- dictionary on page 1043
- request services extensible-subscriber-services reload-dictionary on page 1536
- show extensible-subscriber-services dictionary on page 1712
- show extensible-subscriber-services dictionary services on page 1719
- Understanding the Dictionary File on page 226

Output Fields Table 115 on page 1716lists the output fields for the **show extensible-subscriber-services dictionary attributes** command. Output fields are listed in the approximate order in which they appear.

Table 115: show extensible-subscriber-services dictionary attributes Output Fields

Field Name	Field Description
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**
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.

Copyright © 2017, Juniper Networks, Inc.

Table 115: show extensible-subscriber-services dictionary attributes Output Fields *(continued)*

Field Name	Field Description
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	Traffic mirroring action. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26–59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26–58, 26–59, 26–60, or 26–61) must always include all four VSAs. If the CoA action is to stop mirroring (VSA 26–58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary attributes

RADIUS DICTIONARY

  Attribute Name              Vendor ID   Attribute Code   Attribute Type  Has Tag  Config Type
  Class                       0           25               string          FALSE    INVALID

  Acct-Session-Id             0           44               string          FALSE    INVALID

  ERX-Service-Activate        4874        65               string          TRUE     SERVICE_ACTIVATE

  Service-Type                4874        173              integer         TRUE     INVALID

  ADSL-Agent-Remote-Id        3561        2                string          FALSE    INVALID

  ERX-Med-Port-Number         4874        61               OctetString     FALSE    INVALID

  NAS-Port-Id                 0           87               text            FALSE    INVALID

  ERX-Med-Dev-Handle          4874        59               OctetString     FALSE    INVALID

  ERX-Service-Deactivate      4874        66               string          TRUE     SERVICE_DEACTIVATE

  NAS-IP-Address              0           4                address         FALSE    INVALID

  NAS-Identifier              0           32               string          FALSE    INVALID
```

Copyright © 2017, Juniper Networks, Inc.

ERX-Service-Acct-Interval	4874	140	integer	TRUE	INVALID
ADSL-Agent-Circuit-Id	3561	1	string	FALSE	INVALID
ERX-LI-Action	4874	58	OctetString	FALSE	SERVICE_ACTIVATE
ERX-Med-Ip-Address	4874	60	OctetString	FALSE	INVALID

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services dictionary services

Syntax show extensible-subscriber-services dictionary services

Release Information Command introduced in Junos OS Release 15.1.

Description Display services configured in the dictionary file and their attributes.

Required Privilege Level view

Related Documentation

- dictionary on page 1043
- request services extensible-subscriber-services reload-dictionary on page 1536
- show extensible-subscriber-services dictionary on page 1712
- show extensible-subscriber-services dictionary attributes on page 1716
- Understanding the Dictionary File on page 226

Output Fields Table 116 on page 1719lists the output fields for the **show extensible-subscriber-services dictionary services** command. Output fields are listed in the approximate order in which they appear.

Table 116: show extensible-subscriber-services dictionary services Output Fields

Field Name	Field Description
Service Name	Name of the service specified in the dictionary.
Service Attribute Name	Name of the service attribute.
PROVISION ACTION	Name, type, and version of the provisioning action
DE-PROVISION ACTION	Name, type, and version of the deprovisioning action
PARAMETERS	Parameters to be processed for the provisioning or deprovisioning action.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary services

RADIUS DICTIONARY

  Services List

  Service Name: ngcoco
    Service Attribute Name      : ERX-Service-Activate
```

Copyright © 2017, Juniper Networks, Inc.

```
    PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_ngcoco_add

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

  Service Name: dhcprelay
    Service Attribute Name      : ERX-Service-Activate

    PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_dhcprelay_add

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate
  Service Name: default
    Service Attribute Name      : ERX-LI-Action

    PROVISION ACTION
      Action Type               : APPLICATION
      Action Version            : 1
      Action Name               : LI

      PARAMETERS
        ERX-Med-Dev-Handle
        ERX-Med-Ip-Address
        ERX-Med-Port-Number

  Service Name: dhcprelay
    Service Attribute Name      : ERX-Service-Deactivate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
```


Copyright © 2017, Juniper Networks, Inc.

```
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Deactivate

Service Name: ngcoco
  Service Attribute Name      : ERX-Service-Deactivate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Deactivate
```

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services sessions

Syntax show extensible-subscriber-services sessions *accounting-session-id*

Release Information Command introduced in Junos OS Release 15.1.

Description Display session information. It displays a list of the services applied and their current state. Specify the accounting session ID to view information about a specific session. If an accounting session ID is not specified in the command, the command displays details of all the sessions.

Options ***accounting-session-id***—Identifier of the session you want information about.

Required Privilege Level view

Related Documentation
- clear extensible-subscriber-services sessions on page 1503

Output Fields Table 117 on page 1722 describes the output fields for the **show extensible-subscriber-services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 117: show extensible-subscriber-services sessions Output Fields

Field Name	Field Description
Session ID	Session ID. Accounting session ID must be enclosed in quotation marks.
Timestamp	Time when the session was started
Service Name	Name of the service
Commit State	Commit state: **Init**, **Queued**, or **Success**.
Service ID	Service ID

Sample Output

```
#show extensible-subscriber-services sessions
Session ID: jnpr demux0.1073762028:46422
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Service Name: dhcprelay

Total Sessions: 1

# show extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"

Service ID: jnpr demux0.1073762028:46422:46425
```


Copyright © 2017, Juniper Networks, Inc.

```
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Commit State: Success

Service ID: jnpr demux0.1073762028:46422:46426
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: dhcprelay
    Commit State: Success

Total Services: 2
```

Copyright © 2017, Juniper Networks, Inc.

show extensible-subscriber-services service

Syntax show extensible-subscriber-services service

Release Information Command introduced in Junos OS Release 15.1.

Description Display service information. It displays the service name specified in the dictionary, scripts used for provisioning and deprovisioning, and the number of services. A service can have more than one version. Each version is grouped using an action set. The number of action sets is equal to the number of versions for the service. Number of services in all the versions is displayed in **Total Services**.

Required Privilege Level view

Related Documentation
- Understanding the Dictionary File on page 226

Output Fields Table 118 on page 1724 lists the output fields for the **show extensible-subscriber-services service** command. Output fields are listed in the approximate order in which they appear.

Table 118: show extensible-subscriber-services service Output Fields

Field Name	Field Description
Service Name	Name of the service
Action	Name and number of services provisioned or deprovisioned.
Total Services	Total number of services provisioned or deprovisioned.

Sample Output

```
#show extensible-subscriber-services services
Service Name: ngcoco
    Action:
        Provision:         iceaaa_ngcoco_add_1
        Deprovision:       iceaaa_del_1
        Services:           1301

    Total Services:           1301

Service Name: dhcprelay
    Action:
        Provision:         iceaaa_dhcprelay_add_1
        Deprovision:       iceaaa_del_1
        Services:           1301

    Total Services:          1301
```


Copyright © 2017, Juniper Networks, Inc.

show ipv6 router-advertisement

Syntax

show ipv6 router-advertisement
<conflicts>
<interface *interface*>
<logical-system (all | *logical-system-name*)>
<prefix *prefix/prefix length*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

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

Related Documentation

- clear ipv6 router-advertisement on page 1504

List of Sample Output

show ipv6 router-advertisement on page 1726
show ipv6 router-advertisement conflicts on page 1727
show ipv6 router-advertisement prefix on page 1727

Output Fields

Table 119 on page 1726 describes the output fields for the **show ipv6 router-advertisement** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 119: 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.
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).

Sample Output

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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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 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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access aaa accounting

Syntax **show network-access aaa accounting**

Release Information Command introduced in Junos OS Release 12.3.

Description Display the state of the RADIUS Acct-On response sent from the RADIUS server.

Required Privilege Level view

Related Documentation
- RADIUS Acct-On and Acct-Off Messages on page 105

List of Sample Output show network-access aaa accounting on page 1728

Output Fields Table 120 on page 1728 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 120: 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.
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
```


Copyright © 2017, Juniper Networks, Inc.

show network-access aaa radius-servers

Syntax **show network-access aaa radius-servers**
<detail>

Release Information Command introduced in Junos OS Release 12.1.

Description Display RADIUS server status and information.

Options **detail**—(Optional) Display detailed level of information.

Required Privilege Level view

Related Documentation
- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output
show network-access aaa radius-servers on page 1733
show network-access aaa radius-servers on page 1734
show network-access aaa radius-servers detail on page 1734

Output Fields Table 121 on page 1729 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 121: 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

Copyright © 2017, Juniper Networks, Inc.

Table 121: 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**. 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. For example, if the subsequent subscriber logs in a minute after the request and all retries lapse, and the 10-second timeout starts, the actual time until the server is marked **Down** is 50 seconds after the timeout starts (the one minute between subscriber login minus the 10-second timeout). • 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**. • 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**.	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 121: show network-access aaa radius-servers Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 121: show network-access aaa radius-servers Output Fields *(continued)*

Field Name	Field Description	Level of Output
Preauthentication retransmissions	Number of retransmissions of preauthentication requests.	Detail
Preauthentication Accepts	Number of preauthentication requests accepted by the preauthentication server.	Detail
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 121: show network-access aaa radius-servers Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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
```

Copyright © 2017, Juniper Networks, Inc.

```
      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
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Accounting retry: 5
      Accounting port: 60
      Status: UP
    Server address: 192.168.30.192
      Authentication port: 1812
      Preauthentication port: 1810
      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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access aaa statistics

Syntax

show network-access aaa statistics
<accounting (detail)>
<address-assignment (client | pool *pool-name*)>
<dynamic-requests>
<radius>

Release Information

Command introduced in Junos OS Release 9.1.
Option **address-assignment** introduced in Junos OS Release 10.0.
Option **radius** introduced in Junos OS Release 11.4.
Option **detail** introduced in Junos OS Release 13.3.

Description

Display AAA accounting, address-assignment, dynamic request statistics, and RADIUS settings and 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.

Required Privilege Level

view

Related Documentation

- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output

show network-access aaa statistics accounting on page 1742
show network-access aaa statistics accounting detail on page 1743
show network-access aaa statistics address-assignment client on page 1743
show network-access aaa statistics address-assignment pool on page 1743
show network-access aaa statistics dynamic-requests on page 1744
show network-access aaa statistics radius on page 1744

Output Fields

Table 122 on page 1737 lists the output fields for the **show network-access aaa statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 122: 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**
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

Copyright © 2017, Juniper Networks, Inc.

Table 122: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 122: 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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 122: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**
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**

Copyright © 2017, Juniper Networks, Inc.

Table 122: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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
Address Usage (percent)	Percentage of total addresses in use.	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
Link Name	Name of the secondary address-assignment pool to which the primary pool is linked.	

Copyright © 2017, Juniper Networks, Inc.

Table 122: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

```
Out of Addresses: 9
Address total: 47
Addresses in use: 47
Address Usage (percent): 100
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
```


Copyright © 2017, Juniper Networks, Inc.

show network-access aaa statistics authentication

Syntax **show network-access aaa statistics authentication <detail>**

Release Information Command introduced in Junos OS Release 9.1.
Option **detail** introduced in Junos OS Release 12.1.

Description Display AAA authentication statistics.

Options **detail**—(Optional) Displays detailed information about authentication.

Required Privilege Level view

Related Documentation
- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output
show network-access aaa statistics authentication on page 1747
show network-access aaa statistics authentication detail on page 1747

Output Fields Table 123 on page 1745 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 123: 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
Queue request deleted	Number of queue requests that have been deleted.	Detail

Copyright © 2017, Juniper Networks, Inc.

Table 123: show network-access aaa statistics authentication Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access aaa statistics pending-accounting-stops

Syntax **show network-access aaa statistics pending-accounting-stops**

Release Information Command introduced in Junos OS Release 13.1.

Description Display the number of pending accounting stop requests.

Options This command has no options.

Required Privilege Level view

Related Documentation
- request network-access aaa replay pending-accounting-stops on page 1334
- show accounting pending-accounting-stops on page 1548

List of Sample Output show network-access aaa statistics pending-accounting-stops on page 1748

Output Fields Table 124 on page 1748 lists the output field for the **show network-access aaa statistics pending-accounting-stops** command.

Table 124: 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
```


Copyright © 2017, Juniper Networks, Inc.

show network-access aaa statistics preauthentication

Syntax **show network-access aaa statistics preauthentication**

Release Information Command introduced in Junos OS Release 13.3.

Description Display AAA preauthentication statistics.

Options **detail**—(Optional) Displays detailed information about authentication.

Required Privilege Level view

Related Documentation
- RADIUS Logical Line Identifier (LLID) Overview on page 149
- Configuring Logical Line Identification (LLID) Preauthentication on page 152

List of Sample Output show network-access aaa statistics preauthentication on page 1750

Output Fields Table 125 on page 1749 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 125: 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
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

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

show network-access aaa subscribers

Syntax

show network-access aaa subscribers
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>
<statistics>
<username>

Release Information

Command introduced in Junos OS Release 9.1.

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.

Required Privilege Level

view

Related Documentation

- Verifying and Managing Subscriber AAA Information on page 227

List of Sample Output

show network-access aaa subscribers logical-system on page 1752
show network-access aaa subscribers logical-system routing-instance on page 1752
show network-access aaa subscribers statistics username on page 1753
show network-access aaa subscribers username on page 1753

Output Fields

Table 126 on page 1751 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 126: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 126: show network-access aaa subscribers Output Fields *(continued)*

Field Name	Field Description
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.
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 aborted 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.

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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access aaa subscribers session-id

Syntax **show network-access aaa subscribers session-id** *session-id*
<brief | detail>

Release Information Command introduced in Junos OS Release 10.0.

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

Related Documentation

- Verifying and Managing Subscriber AAA Information on page 227
- Activating and Deactivating Subscriber Services Locally with the CLI on page 610
- Deactivating a Single Instance of a Subscriber Service with Multiple Instances on page 619
- Deactivating All Instances of a Subscriber Service with Multiple Instances on page 622
- Verifying and Managing Subscriber Services with Multiple Instances on page 625
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

List of Sample Output

show network-access aaa subscribers session-id brief on page 1758
show network-access aaa subscribers session-id 2 (Flat File Accounting) on page 1758
show network-access aaa subscribers session-id detail on page 1759
show network-access aaa subscribers session-id detail (Service with Multiple Instances) on page 1759
show network-access aaa subscribers session-id detail (Single Session Dual Stack with active V4 and V6 subscribers) on page 1760

Output Fields Table 127 on page 1755 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.

Copyright © 2017, Juniper Networks, Inc.

Table 127: 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
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

Copyright © 2017, Juniper Networks, Inc.

Table 127: show network-access aaa subscribers session-id Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 127: show network-access aaa subscribers session-id Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Copyright © 2017, Juniper Networks, Inc.

Table 127: 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

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
```


Copyright © 2017, Juniper Networks, Inc.

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 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: 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
Accounting State: Acc-Interim-Sent
```

Copyright © 2017, Juniper Networks, Inc.

```
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  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: retal-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: 3000:0:0:8003::2
IPv6 Prefix: 3ffe:ffff:0:4::/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```


Copyright © 2017, Juniper Networks, Inc.

show network-access aaa terminate-code

Syntax

show network-access aaa terminate-code
<brief | detail | summary>
<reverse>
<(aaa | dhcp | l2tp | ppp)>

Release Information

Command introduced in Junos OS Release 11.4.
vlan option introduced in Junos OS Release 16.1.

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

Related Documentation

- Understanding Session Termination Causes and RADIUS Termination Cause Codes on page 187
- Mapping Session Termination Causes to Custom Termination Cause Codes on page 189
- AAA Termination Causes and Code Values on page 190
- DHCP Termination Causes and Code Values on page 192
- L2TP Termination Causes and Code Values on page 193
- PPP Termination Causes and Code Values on page 213
- VLAN Termination Causes and Code Values on page 221

Copyright © 2017, Juniper Networks, Inc.

List of Sample Output

show network-access aaa terminate-code on page 1762
show network-access aaa terminate-code reverse on page 1763
show network-access aaa terminate-code dhcp on page 1763
show network-access aaa terminate-code detail on page 1763
show network-access aaa terminate-code brief on page 1764
show network-access aaa terminate-code summary on page 1764
show network-access aaa terminate-code vlan on page 1764
show network-access aaa terminate-code vlan detail on page 1764

Output Fields

Table 128 on page 1762 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 128: 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
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
```


Copyright © 2017, Juniper Networks, Inc.

```
  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:
```

Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access address-assignment pool

Syntax show network-access address-assignment pool *pool-name*
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 9.0.

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

List of Sample Output show network-access address-assignment pool on page 1766

Output Fields Table 129 on page 1766 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 129: show network-access address-assignment pool Output Fields

Field Name	Field Description
IP address	IP address of the client.
Hardware address	MAC address of the client.
Type	Type of client.

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       Hardware address       Type
192.168.2.1      00:00:5e:00:53:01      DHCP
192.168.2.2      00:00:5e:00:53:02      DHCP
```


Copyright © 2017, Juniper Networks, Inc.

```
192.168.2.3       00:00:5e:00:53:03       DHCP
192.168.2.4       00:00:5e:00:53:04       DHCP
```

Copyright © 2017, Juniper Networks, Inc.

show network-access domain-map

Syntax **show network-access domain-map**
<statistics>

Release Information Command introduced in Junos OS Release 10.4.

Description Display domain map information.

Options **statistics**—(Optional) Display domain map statistics.

Required Privilege Level view

Related Documentation
- Verifying and Managing Domain Map Configuration on page 230

List of Sample Output show network-access domain-map statistics on page 1768

Output Fields Table 130 on page 1768 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 130: 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
```


Copyright © 2017, Juniper Networks, Inc.

```
Domain statistics for domain-name: default
  Default used: 1
```

Copyright © 2017, Juniper Networks, Inc.

show network-access gx-plus

Syntax **show network-access gx-plus**
<state | statistics | sync-state>

Release Information Command introduced in Junos OS Release 11.2.

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

Related Documentation
- clear network-access gx-plus replay on page 1509
- clear network-access gx-plus statistics on page 1510

List of Sample Output
show network-access gx-plus state on page 1771
show network-access gx-plus statistics on page 1771
show network-access gx-plus sync-state on page 1771

Output Fields Table 131 on page 1770 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 131: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 131: show network-access gx-plus Output Fields *(continued)*

Field Name	Field Description
Gx-plus general counters	Number and state of general events.
Gx-plus sync-event counters	Number and state of synchronization events.

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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access nasreq statistics

Syntax show network-access nasreq statistics

Release Information Command introduced in Junos OS Release 16.1.

Description Display information about the NASREQ protocol exchanges that are used to authenticate and authorize subscribers when they log in.

Required Privilege Level view

Related Documentation

- Diameter Network Access Server Application (NASREQ) on page 779
- Configuring the Diameter Network Access Server Application (NASREQ) on page 781
- Messages Used by Diameter Applications on page 706

Output Fields Table 132 on page 1772 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 132: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 132: show network-access nasreq statistics Output Fields *(continued)*

Field Name	Field Description
Authorization request messages timeouts	Number of authorization requests that timed out waiting for a response from the NASREQ server.
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 requests received by the NASREQ engine.
Abort-session response-ack messages sent	Number of abort-session requests receiving an ACK response.
Abort-session response-nack messages sent	Number of abort-session requests receiving a NACK response.
Abort-session response message failures	Number of abort-session responses 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.

Copyright © 2017, Juniper Networks, Inc.

Table 132: show network-access nasreq statistics Output Fields *(continued)*

Field Name	Field Description
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.
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

```
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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access ocs

Syntax

show network-access ocs
<state>
<statistics [brief | detail | extensive]>

Release Information

Command introduced in Junos OS Release 16.2.

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

Related Documentation

- ocs on page 1218
- clear network-access ocs statistics on page 1511
- Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers on page 744
- Understanding Gy Interactions Between the Router and the OCS on page 756
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761

List of Sample Output

show network-access ocs state on page 1777
show network-access ocs statistics on page 1777

Output Fields

Table 133 on page 1776 lists the output fields for the **show network-access ocs** command. Output fields are listed in the approximate order in which they appear.

Table 133: 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

Copyright © 2017, Juniper Networks, Inc.

Table 133: show network-access ocs Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```

Copyright © 2017, Juniper Networks, Inc.

show network-access pcrf

Syntax

show network-access pcrf
<state>
<statistics [brief | detail | extensive]>

Release Information

Command introduced in Junos OS Release 16.2.

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.

Required Privilege Level

view

Related Documentation

- pcrf on page 1259
- clear network-access pcrf on page 1512
- Understanding Gx Interactions Between the Router and the PCRF on page 748
- Understanding Rule Interactions Between the PCRF, PCEF, and OCS on page 761
- Understanding Upstream and Downstream Messages for the PCRF on page 766

List of Sample Output

show network-access pcrf state on page 1779
show network-access pcrf statistics on page 1779

Output Fields

Table 134 on page 1779 lists the output fields for the **show network-access pcrf** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 134: 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. • **started**—Log in. • **local-active**—Services installed after local grant. • **local-deny**—Denied by local decision waiting for subscriber logout. • **local-grant**—Recieved local grant and installing services. • **grant**—Received remote grant and installing services. • **upd-from-local**—Received remote grant, in local-active state, and installing services. • **update**—Processing RAR-update. • **report**—Sending report. • **deny**—Received deny from remote and waiting for subscriber logout. • **failed**—GRES before active state or provisioning failed, waiting for subscriber logout. • **ocs-final-wait**—Waiting for OCS to complete final interrogation. • **logout**—Waiting for logout response. • **dead**—Received logout response or timed out waiting for logout response.	All levels
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

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
```

Copyright © 2017, Juniper Networks, Inc.

```
  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
```


Copyright © 2017, Juniper Networks, Inc.

show ppp address-pool

Syntax show ppp address-pool *pool-name*
<detail>

Release Information Command introduced in Junos OS Release 7.5.

Description Display PPP address pool information.

Options ***pool-name***—Address pool name.

detail—(Optional) Display detailed address pool information.

Required Privilege Level view

Related Documentation
- *Verifying and Managing PPP Configuration for Subscriber Management*

List of Sample Output
show ppp address-pool on page 1782
show ppp address-pool detail on page 1782

Output Fields Table 135 on page 1781 lists the output fields for the **show ppp address-pool** command. Output fields are listed in the approximate order in which they appear.

Table 135: 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**

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

show route

List of Syntax

Syntax on page 1783
Syntax (EX Series Switches) on page 1783

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>
<te-ipv4-prefix-ip *te-ipv4-prefix-ip*>
<te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip*>
<te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso*>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.
Options **te-ipv4-prefix-ip**, **te-ipv4-prefix-node-ip**, and **te-ipv4-prefix-node-iso** introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

display-client-data —(Optional) Display client id and cookie information for routes installed by rpd client applications.

te-ipv4-prefix-ip ***te-ipv4-prefix-ip***—(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.

Copyright © 2017, Juniper Networks, Inc.

te-ipv4-prefix-node-ip ***te-ipv4-prefix-node-ip***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the **lsdist.0** table.

te-ipv4-prefix-node-iso ***te-ipv4-prefix-node-iso***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the **lsdist.0** table.

Required Privilege Level view

Related Documentation

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

List of Sample Output

show route on page 1787
show route (VPN) on page 1788
show route (with Destination Prefix) on page 1788
show route destination-prefix detail on page 1788
show route extensive on page 1788
show route extensive (ECMP) on page 1789
show route (Enhanced Subscriber Management) on page 1789
show route (IPv6 Flow Specification) on page 1789
show route display-client-data detail on page 1790
show route te-ipv4-prefix-ip on page 1790
show route te-ipv4-prefix-ip extensive on page 1791
show route te-ipv4-prefix-node-iso on page 1794
show route te-ipv4-prefix-node-iso extensive on page 1794
show route te-ipv4-prefix-node-iso detail on page 1797

Output Fields Table 136 on page 1784 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 136: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 136: show route Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).

destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96):
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Copyright © 2017, Juniper Networks, Inc.

Table 136: show route Output Fields *(continued)*

Field Name	Field Description
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
Route Labels	Stack of labels carried in the BGP route update.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.

Copyright © 2017, Juniper Networks, Inc.

Table 136: show route Output Fields *(continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```

Copyright © 2017, Juniper Networks, Inc.

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 192.0.2.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                   #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 198.51.100.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
```


Copyright © 2017, Juniper Networks, Inc.

```
            PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group
203.0.113.1
            Next hop type: Indirect
            Address: 0x92455b8
            Next-hop reference count: 2
            Source: 10.0.0.30
            Protocol next hop: 10.0.0.40
            Indirect next hop: 2 no-forward
            State: <Active Int Ext>
                     Local AS: 64510 Peer AS: 64511
            Age: 3  Metric2: 1
           Validation State: unverified
            Task: BGP_64510.10.0.0.30+179
            Announcement bits (2): 0-PIM.v1 1-mvpn global task
            AS path: I (Originator) Cluster list:  10.0.0.30
            AS path:  Originator ID: 10.0.0.40
            Communities: target:64502:100
            Import Accepted
            Localpref: 100
            Router ID: 10.0.0.30
            Primary Routing Table bgp.mvpn.0
            Indirect next hops: 1
                 Protocol next hop: 10.0.0.40 Metric: 1
                 Indirect next hop: 2 no-forward
                 Indirect path forwarding next hops: 1
                       Next hop type: Router
                       Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                 10.0.0.40/32 Originating RIB: inet.3
                       Metric: 1                       Node path count: 1
                       Forwarding nexthops: 1
                          Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route extensive (ECMP)

```
user@host> show route extensive
*IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 172.16.50.2 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 192.0.2.2 via ae0.0 balance 57%
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                    > to #0 10.0.0.1.93.65 via demux0.1073741824
198.51.100.12/24    *[Access-internal/12] 00:00:08
                    Private unicast
```

show route (IPv6 Flow Specification)

```
user@host> show route
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

Copyright © 2017, Juniper Networks, Inc.

```
2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

show route display-client-data detail

```
user@host> show route 198.51.100.0/24 display-client-data detail
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
        *BGP-Static Preference: 5/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xa5c2af8
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 1641
                Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                Session Id: 0x160
                Protocol next hop: 192.0.2.1
                Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                Age: 3:13       Metric2: 0
                Validation State: unverified
                Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
                AS path: I
                Client id: 1, Cookie: 1
```

show route te-ipv4-prefix-ip

```
user@host> show route te-ipv4-prefix-ip 10.10.10.10
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


Copyright © 2017, Juniper Networks, Inc.

```
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                    *[IS-IS/15] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
```

show route te-ipv4-prefix-ip extensive

```
user@host>show route te-ipv4-prefix-ip 10.10.10.10 extensive
lsdist.0: 298 destinations, 298 routes (298 active, 0 holddown, 0 hidden)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State:<Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
```

Copyright © 2017, Juniper Networks, Inc.

```
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
        State: <Active NotInstall
        Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0>

PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
        Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
```


Copyright © 2017, Juniper Networks, Inc.

```
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

Copyright © 2017, Juniper Networks, Inc.

show route te-ipv4-prefix-node-iso

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
```

show route te-ipv4-prefix-node-iso extensive

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 extensive
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
```

Copyright © 2017, Juniper Networks, Inc.

```
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 283
                    Next hop:
                    State: <Active NotInstall>
                    Local AS:   100
                    Age: 6:47
                    Validation State: unverified
                    Task: IS-IS
                    AS path: I
                    Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso detail

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 detail
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

```
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
```

 Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

Copyright © 2017, Juniper Networks, Inc.

show route extensive

List of Syntax
Syntax on page 1800
Syntax (EX Series Switches) on page 1800

Syntax
show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route extensive
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display extensive information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route extensive on page 1807
show route extensive (Access Route) on page 1814
show route extensive (BGP PIC Edge) on page 1814
show route extensive (FRR and LFA) on page 1815
show route extensive (IS-IS) on page 1816
show route extensive (Route Reflector) on page 1816
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1816
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 1817

Output Fields
Table 137 on page 1800 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 137: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

***route-destination* (entry, announced)**	Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the *show route detail* command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the *show route detail* command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown. For sample output, see **show route table**.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see **show route table**.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
route status	Indicates the status of a BGP route: • **Accepted**—The specified BGP route is imported by the default BGP policy. • **Import**—The route is imported into a Layer 3 VPN routing instance. • **Import-Protect**—A remote instance egress that is protected. • **Multipath**—A BGP multipath active route. • **MultipathContrib**—The route is not active but contributes to the BGP multipath. • **Protect**—An egress route that is protected. • **Stale**—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the *show route detail* command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 137: show route extensive Output Fields *(continued)*

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```

Copyright © 2017, Juniper Networks, Inc.

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                               Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
```


Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
```

```
                    Route Label: 300112
                    Localpref: 100
                    Router ID: 213.0.113.99
```

show route extensive (Access Route)

```
user@host> show route 203.0.113.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I


user@host> show route 2001:db8:4641:1::/48 extensive

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 198.51.100.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
```


Copyright © 2017, Juniper Networks, Inc.

```
          State: <CalcForwarding>
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
          #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
                  Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                  Indirect next hop: weight 0x4000
                  State: <ForwardingOnly Int Ext)>
                  Inactive reason: Forwarding use only
                  Age: 25         Metric2: 15
                  Validation State: unverified
                  Task: RT
                  Announcement bits (1): 0-KRT
                  AS path: 3 I
                  Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 203.0.113.20 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0xb9d78c4
                    Next-hop reference count: 7
                    Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                    Session Id: 0x201
                    State: Int
                    Inactive reason: Route Preference
                    Local AS:   64500
                    Age: 5:35  Metric: 3
                    Area: 0.0.0.0
                    Task: OSPF
                    AS path: I
```

show route extensive (IS-IS)

```
user@host> show route extensive
IS-IS Preference: 15
                    Level: 1
                    Next hop type: Router, Next hop index: 1048577
                    Address: 0xXXXXXXXXXX
                    Next-hop reference count: YY
                    Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                    Session Id: 0x141
                    Next hop: 203.0.113.22 via ae0.0 balance 57%
```

show route extensive (Route Reflector)

```
user@host> show route extensive
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 203.0.113.4
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: via vt-0/1/0.1
                    Next-hop index: 661
                    Label operation: Pop
                    Address: 0x9172130
                    Next hop: via so-0/0/3.0
                    Next-hop index: 654
                    Label operation: Swap 299872
                    State: **Active Int>
                    Local AS:  64511
                    Age: 8:20       Metric: 1
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 203.0.113.82 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 203.0.113.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    64500
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 198.51.100.1, grp: 232.1.1.1,
 src: 192.168.219.11
                Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                   RPF Nexthops :
                       ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                       ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```

Copyright © 2017, Juniper Networks, Inc.

show static-subscribers sessions

Syntax **show static-subscribers sessions**
<group *group-name*>
<interface *interface-name*>

Release Information Command introduced in Junos OS Release 9.6.

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

Related Documentation
- Subscribers on Static Interfaces Overview on page 797

List of Sample Output
show static-subscribers sessions on page 1819
show static-subscribers sessions group on page 1819
show static-subscribers sessions interface on page 1819

Output Fields Table 138 on page 1818 lists the output fields for the **show static-subscribers sessions** command. Output fields are listed in the approximate order in which they appear.

Table 138: show static-subscribers sessions Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	None specified
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

Copyright © 2017, Juniper Networks, Inc.

Table 138: show static-subscribers sessions Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```

Copyright © 2017, Juniper Networks, Inc.

show subscribers

Syntax

show subscribers
<detail | extensive | terse>
<aci-interface-set-name *aci-interface-set-name*>
<address *address*>
<agent-circuit-identifier *agent-circuit-identifier-substring*>
<client-type *client-type*>
<count>
<id>
<interface *interface*>
<logical-system *logical-system*>
<mac-address *mac-address*>
<physical-interface *physical-interface-name*>
<profile-name *profile-name*>
<routing-instance *routing-instance*>
<stacked-vlan-id *stacked-vlan-id*>
<subscriber-state *subscriber-state*>
<user-name *user-name*>
<vci *vci-identifier*>
<vpi *vpi-identifier*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 9.3.
Command introduced in Junos OS Release 9.3 for EX Series switches.
client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.
Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Display information for active subscribers.

Options

detail | extensive | terse—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

Copyright © 2017, Juniper Networks, Inc.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dotlx**—Dotlx clients only.
- **essm**—ESSM clients only.
- **fwauth**—FwAuth (authenticated across a firewall) clients only.
- **l2tp**—L2TP clients only.
- **mlppp**—MLPPP clients only.
- **ppp**—PPP clients only.
- **pppoe**—PPPoE clients only.
- **static**—Static clients only.
- **vlan**—VLAN clients only.
- **vlan-oob**—VLAN out-of-band (ANCP-triggered) clients only.
- **vpls-pw**—VPLS pseudowire clients only.
- **xauth**—Xauth clients only.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified 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.

Copyright © 2017, Juniper Networks, Inc.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** ***stacked-vlan-id*** option to match the outer VLAN tag.

NOTE: **Because of display limitations, logical system and routing instance output values are truncated when necessary.**

Required Privilege Level view

Related Documentation

- show subscribers summary on page 1846
- *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 on page 577

List of Sample Output

show subscribers (IPv4) on page 1829
show subscribers (IPv6) on page 1829

Copyright © 2017, Juniper Networks, Inc.

show subscribers (IPv4 and IPv6 Dual Stack) on page 1829
show subscribers (Single Session DHCP Dual Stack) on page 1829
show subscribers (Single Session DHCP Dual Stack detail) on page 1830
show subscribers (LNS on MX Series Routers) on page 1830
show subscribers (L2TP Switched Tunnels) on page 1830
show subscribers client-type dhcp detail on page 1830
show subscribers client-type vlan-oob detail on page 1831
show subscribers count on page 1831
show subscribers address detail (IPv6) on page 1831
show subscribers detail (IPv4) on page 1832
show subscribers detail (IPv6) on page 1832
show subscribers detail (pseudowire Interface for GRE Tunnel) on page 1832
show subscribers detail (IPv6 Static Demux Interface) on page 1833
show subscribers detail (L2TP LNS Subscribers on MX Series Routers) on page 1833
show subscribers detail (L2TP Switched Tunnels) on page 1833
show subscribers detail (Tunneled Subscriber) on page 1834
show subscribers detail (IPv4 and IPv6 Dual Stack) on page 1834
show subscribers detail (ACI Interface Set Session) on page 1835
show subscribers detail (PPPoE Subscriber Session with ACI Interface Set) on page 1835
show subscribers extensive on page 1835
show subscribers extensive (Passive Optical Network Circuit Interface Set) on page 1836
show subscribers extensive (DNS Addresses from Access Profile or Global Configuration) on page 1836
show subscribers extensive (DNS Addresses from RADIUS) on page 1837
show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration) on page 1837
show subscribers extensive (RPF Check Fail Filter) on page 1838
show subscribers extensive (L2TP LNS Subscribers on MX Series Routers) on page 1838
show subscribers extensive (IPv4 and IPv6 Dual Stack) on page 1838
show subscribers extensive (ADF Rules) on page 1839
show subscribers extensive (Effective Shaping-Rate) on page 1840
show subscribers extensive (Subscriber Session Using PCEF Profile) on page 1840
show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set) on page 1841
show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring) on page 1842
show subscribers interface extensive on page 1842
show subscribers logical-system terse on page 1843
show subscribers physical-interface count on page 1843
show subscribers routing-instance inst1 count on page 1843
show subscribers stacked-vlan-id detail on page 1843
show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 1843
show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 1844
show subscribers user-name detail on page 1844
show subscribers vlan-id on page 1844
show subscribers vlan-id detail on page 1844
show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session) on page 1845
show subscribers address detail (Enhanced Subscriber Management) on page 1845

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 139 on page 1824 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 139: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is **Tunnel-switched**.
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask. (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
Primary DNS Address	IP address of primary DNS server. This field is displayed with the **extensive** option only when the address is provided by RADIUS.
Secondary DNS Address	IP address of secondary DNS server. This field is displayed with the **extensive** option only when the address is provided by RADIUS.
IPv6 Primary DNS Address	IPv6 address of primary DNS server. This field is displayed with the **extensive** option only when the address is provided by RADIUS.
IPv6 Secondary DNS Address	IPv6 address of secondary DNS server. This field is displayed with the **extensive** option only when the address is provided by RADIUS.
Domain name server inet	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the **extensive** option only when the addresses are derived from the access profile or the global access configuration.

Copyright © 2017, Juniper Networks, Inc.

Table 139: show subscribers Output Fields *(continued)*

Field Name	Field Description
Domain name server inet6	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the **extensive** option only when the addresses are derived from the access profile or the global access configuration.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form **demux0.*nnnn*** (for example, **demux0.3221225472**), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is **Static** or **Dynamic**.

Copyright © 2017, Juniper Networks, Inc.

Table 139: show subscribers Output Fields *(continued)*

Field Name	Field Description
Interface Set	Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line. The name of the interface set uses one of the following prefixes: • **aci**—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • **ari**—ARI; for example, ari-1033-demux0.3221225524. • **aci+ari**—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524. • **noids**—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. NOTE: ACI interface sets are configured with the **agent-circuit-identifier** autoconfiguration stanza. ALI interface sets are configured with the **line-identity** autoconfiguration stanza. Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable $junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.
Interface Set Type	Interface type of the ACI interface set: **Dynamic**. This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (**Init**, **Configured**, **Active**, **Terminating**, **Tunneled**).
L2TP State	Current state of the L2TP session, **Tunneled** or **Tunnel-switched**. When the value is **Tunnel-switched**, two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.

Copyright © 2017, Juniper Networks, Inc.

Table 139: show subscribers Output Fields *(continued)*

Field Name	Field Description
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	For the **dhcp** client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the **vlan-oob** client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.
Agent Remote ID	For the **dhcp** client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the **vlan-oob** client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.
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.
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.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.

Copyright © 2017, Juniper Networks, Inc.

Table 139: show subscribers Output Fields *(continued)*

Field Name	Field Description
DHCP Options	**len** = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable $junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).

Copyright © 2017, Juniper Networks, Inc.

Table 139: show subscribers Output Fields *(continued)*

Field Name	Field Description
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).

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
```

Copyright © 2017, Juniper Networks, Inc.

show subscribers (Single Session DHCP Dual Stack detail)

```
user@host> show subscribers id 27 detail
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00
```

show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name         LS:RI
si-4/0/0.1         192.0.2.0          user@example.com   default:default
```

show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface           IP Address/VLAN ID    User Name              LS:RI
si-2/1/0.1073741842 Tunnel-switched      user@example.com       default:default

si-2/1/0.1073741843 Tunnel-switched      user@example.com       default:default
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT
```


Copyright © 2017, Juniper Networks, Inc.

```
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 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
```

Copyright © 2017, Juniper Networks, Inc.

```
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      30.1.0.2
ps0.3221225485      30.1.0.3
demux0.3221225486    1                                        default:default
```


Copyright © 2017, Juniper Networks, Inc.

```
demux0.3221225487    1                                        default:default

demux0.3221225488   100.16.0.1                                default:default

demux0.3221225489   100.16.0.2                                default:default
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: user@example.net
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.net
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```


Copyright © 2017, Juniper Networks, Inc.

```
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 extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
```

Copyright © 2017, Juniper Networks, Inc.

```
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (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: 00:10:94:00:00:02
State: Active
Radius Accounting ID: jnpr :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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

```
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user@example.net
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
```


Copyright © 2017, Juniper Networks, Inc.

```
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 (Subscriber Session Using PCEF Profile)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 5.0.0.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:11:01:00:00:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

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 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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

show subscribers summary

Syntax

show subscribers summary
<all>
<detail | extensive | terse>
<count>
<physical-interface *physical-interface-name*>
<logical-system *logical-system* pic | port | routing-instance *routing-instance* | slot>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display summary information for subscribers.

Options

none—Display summary information by state and client type for all subscribers.

all—(Optional) Display summary information by state, client type, and LS:RI.

detail | extensive | terse—(Not supported on MX Series routers) (Optional) Display the specified level of output.

count—(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system *logical-system*—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface *physical-interface-name*—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type, and LS:RI.

pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.

port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.

routing-instance *routing-instance*—(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.

NOTE: **Due to display limitations, logical system and routing instance output values are truncated when necessary.**

Required Privilege Level

view

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- show subscribers on page 1820

List of Sample Output

show subscribers summary on page 1848
show subscribers summary all on page 1848
show subscribers summary physical-interface on page 1849
show subscribers summary physical-interface pic on page 1849
show subscribers summary physical-interface port on page 1849
show subscribers summary physical-interface slot on page 1850
show subscribers summary pic on page 1850
show subscribers summary pic (Aggregated Ethernet Interfaces) on page 1850
show subscribers summary port on page 1850
show subscribers summary port extensive on page 1850
show subscribers summary slot on page 1851
show subscribers summary terse on page 1851

Output Fields

Table 140 on page 1847 lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 140: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
Subscribers by State	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.	**detail** none
Subscribers by Client Type	Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPOE, STATIC-INTERFACE, VLAN, and VLAN-OOB. Also displays the total number of subscribers for all client types (Total).	**detail extensive** none
Subscribers by LS:RI	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).	**detail** none
Subscribers by Connection Type	Number of subscribers summarized by connection type, **Cross-connected** or **Terminated**.	**extensive**
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session. For aggregated Ethernet interfaces, the output of the **summary (pic | port | slot)** options prefixes the interface name with ae0:.	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 140: show subscribers summary Output Fields *(continued)*

Field Name	Field Description	Level of Output
Count	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the **summary** option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.	**detail extensive** none
Total Subscribers	Total number of subscribers for all physical interfaces, all PICS, all ports, or all LS:RI slots.	**detail extensive** none
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*	**terse**
User Name	Name of subscriber.	**terse**
LS:RI	Logical system and routing instance associated with the subscriber.	**terse**

Sample Output

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
 Init      3
 Configured    2
 Active   183
 Terminating    2
 Terminated    1

 TOTAL        191

Subscribers by Client Type
 DHCP     107
 PPP       76
 VLAN         8
 VLAN-OOB       2
 TOTAL        193
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
 Init         3
 Configured       2
 Active      183
 Terminating       2
 Terminated       1

 TOTAL          191

Subscribers by Client Type
 DHCP      107
 PPP          76
```


Copyright © 2017, Juniper Networks, Inc.

```
 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
```

Copyright © 2017, Juniper Networks, Inc.

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 extensive

```
user@host>show subscribers summary port extensive
Interface: ge-5/0/1
Count: 201
Detail:
Subscribers by Client Type
   DHCP: 100
   PPPoE: 100
   VLAN-OOB: 1
Subscribers by Connection Type
   Terminated: 200
   Cross-connected: 1

Interface: ge-5/0/2
Count: 301
Detail:
```


Copyright © 2017, Juniper Networks, Inc.

```
Subscribers by Client Type
   DHCP: 200
   PPPoE: 100
   VLAN-OOB: 1
Subscribers by Connection Type
   Terminated: 300
   Cross-connected: 1

Total Subscribers: 502
```

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface           IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824   100                                   default:default
demux0.1073741824     203.0.113.10         WHOLESALER-CLIENT default:default
demux0.1073741825     203.0.113.13         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     203.0.113.213        RETAILER2-CLIENT  test1:retailer2
```

Copyright © 2017, Juniper Networks, Inc.

show system subscriber-management route

Syntax

show system subscriber-management route
<family *family*>
<incomplete>
<level (brief | detail)>
<next-hop>
<prefix>
<route-type *type*>
<rrt-index *index*>
<summary>

Release Information

Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
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.

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—(Optional) Display the next hop associated with the route entry.

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**.

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.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Verifying and Managing Junos OS Enhanced Subscriber Management on page 577

List of Sample Output
show system subscriber-management route prefix <address> on page 1854
show system subscriber-management route family route-type rtt-index level brief on page 1854
show system subscriber-management route family route-type rtt-index level detail on page 1855
show system subscriber-management route family route-type rtt-index level brief on page 1856

Output Fields Table 141 on page 1853 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 141: 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
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**

Copyright © 2017, Juniper Networks, Inc.

Table 141: show system subscriber-management route Output Fields *(continued)*

Field Name	Field Description	Level of Output
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:               00:00:64:0a:01:0b
     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
```


Copyright © 2017, Juniper Networks, Inc.

```
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
```

Copyright © 2017, Juniper Networks, Inc.

```
    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
2000::/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
```


Copyright © 2017, Juniper Networks, Inc.

show system subscriber-management statistics

Syntax
show system subscriber-management statistics
<all>
<dhcp>
<pppoe>

Release Information
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Enhanced I/O Statistics introduced as part of Extensive output in Junos OS Release 15.1R4 on MX Series routers for enhanced subscriber management.

Description
Display statistics for the specified option. You can customize the output by including one or more optional filters in the command. With the exception of the **extensive** option, all filter options can be combined in a single command.

Options
all —(Optional) Display packet statistics for all protocol.

Required Privilege Level
view

Related Documentation
- *Understanding Dropped Packets and Untransmitted Traffic Using show Commands*

List of Sample Output
show system subscriber-management statistics all on page 1858
show system subscriber-management statistics pppoe on page 1859
show system subscriber-management statistics dhcp on page 1859
show system subscriber-management statistics dhcp extensive on page 1860
show system subscriber-management statistics extensive on page 1862

Output Fields
Table 142 on page 1857 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 142: 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.

Copyright © 2017, Juniper Networks, Inc.

Table 142: show system subscriber-management statistics Output Fields *(continued)*

Field Name	Field Description
ERA discards	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
padis	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
padrs	PPPoE Active Discovery Request packets.
ppp	Point-to-Point Protocol packets.
router solicitations	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.
router advertisements	Number of router advertisements sent or received.
route solicit response packet	Number of router solicitation responses sent or received.

Sample Output

The following examples displays packet statistics accumulated for DHCP and PPPoE since the last time the session manager was cleared.

show system subscriber-management statistics all

```
user@host> show system subscriber-management statistics all
  user@host> show system subscriber-management statistics all
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660

DHCP Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 320008
        ERA discards                    : 6274
    Tx Statistics
        transmit request packets        : 320482
```


Copyright © 2017, Juniper Networks, Inc.

```
        sent packets                    : 320482
    Error Statistics
    Connection Statistics
        no connection packets           : 0

PPPoE Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 486165
        padis                           : 36768
        padrs                           : 35421
        ppp packets                     : 341787
        ERA discards                    : 8249
    Tx Statistics
        packets                         : 70842
        send failures                   : 6240
```

show system subscriber-management statistics pppoe

```
user@host> show system subscriber-management statistics pppoe
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660

PPPoE Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                         : 486165
        padis                           : 36768
        padrs                           : 35421
        ppp packets                     : 341787
        ERA discards                    : 8249
    Tx Statistics
        packets                         : 70842
        send failures                   : 6240
```

show system subscriber-management statistics dhcp

```
user@host> show system subscriber-management statistics dhcp
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
```

```
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
        packets                          : 6604660

DHCP Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                          : 320008
        ERA discards                     : 6274
    Tx Statistics
        transmit request packets         : 320482
        sent packets                     : 320482
    Error Statistics
    Connection Statistics
        no connection packets            : 0
```

show system subscriber-management statistics dhcp extensive

```
user@host> show system subscriber-management statistics dhcp extensive
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
------------------------------------------------------------
                    Packet Statistics
------------------------------------------------------------
I/O Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
    Buffer Statistics
        allocations                      : 7032618
        frees                            : 7032624
        allocation failures              : 0
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
        packets                          : 6604660
    PFE Event Statistics
        packets                          : 0

------------------------------------------------------------
DHCP Statistics:
------------------------------------------------------------
    Rx Statistics
        packets                          : 320008
        ERA discards                     : 6274
    Tx Statistics
        transmit request packets         : 320482
        sent packets                     : 320482
    DHCPv4 Rx Statistics
        total packets                    : 0
    DHCPv4 Tx Statistics
        total packets                    : 0
    DHCPv6 Rx Statistics
```


Copyright © 2017, Juniper Networks, Inc.

```
        total packets                   : 320008
        solicit                         : 36250
        request                         : 36382
        renew                           : 247376
        ERA discards                    : 6274
    DHCPv6 Tx Statistics
        total packets                   : 320482
        advertise                       : 36382
        reply                           : 284100
    Error Statistics
    Connection Statistics
        no connection packets           : 0
        connection down events          : 0
        connection up events            : 0
        flow control invoked            : 0
        flow control released           : 0
        packets sent to daemon          : 320008
        packets received from daemon    : 320482
        messages sent to daemon         : 0
        messages received from daemon   : 0
        notifies while not connected    : 0

NET Statistics:
---------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                        : 36271
        router solicitations            : 36271
    Tx Statistics
        packets:                        : 6284178
        router advertisements           : 6284178
        route solicit response packet   : 36271

Management Statistics:
---------------------------------------------------------------
    dvlan                               : 33912
    dvlan adds                          : 33912
    pppoe                               : 143651
    pppoe add                           : 35750
    pppoe changes                       : 107901
    ip flow                             : 143633
    ip flow add                         : 107883

Management Config Status:
---------------------------------------------------------------
    gres state enabled state            : 1
    shmlog disabled state               : 0
    Rx Statistics
        packets                         : 167361
        ERA discards                    : 15116
    Tx Statistics
        transmit request packets        : 150903
        sent packets                    : 150903
    DHCPv4 Rx Statistics
        total packets                   : 167361
        discover                        : 91910
        request                         : 75451
        ERA discards                    : 15116
```

Copyright © 2017, Juniper Networks, Inc.

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
```


Copyright © 2017, Juniper Networks, Inc.

show system subscriber-management summary

Syntax show system subscriber-management summary

Release Information Command introduced in Junos OS Release 11.1.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Display complete subscriber management database summary information.

Options **none**—This command has no options.

Required Privilege Level view

Related Documentation
- show database-replication statistics on page 1621
- show database-replication summary on page 1623

List of Sample Output show system subscriber-management summary on page 1864
show system subscriber-management summary (Enhanced Subscriber Management) on page 1865

Output Fields Table 143 on page 1863 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 143: 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**

Copyright © 2017, Juniper Networks, Inc.

Table 143: show system subscriber-management summary Output Fields *(continued)*

Field Name	Field Description
Standby	(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine: • **Connected**—Connected but not synchronized • **Disconnected**—Not connected • **Resync (*nn*%)**—Connected and *nn* percent synchronized with the master Routing Engine • **Synchronized**—Synchronized with the master Routing Engine
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
```


Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

test aaa authd-lite user

Syntax test aaa authd-lite user *username* password *password* profile *access-profile-name*
<port *nas-port*>
<zero-stats>

Release Information Command introduced in Junos OS Release 11.2.

Description Verify authd-lite subscriber access authentication, accounting, and address allocation configuration.

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

Related Documentation
- Testing a Subscriber AAA Configuration on page 823

List of Sample Output test aaa authd-lite user on page 1867

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>**.

Copyright © 2017, Juniper Networks, Inc.

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>
```

Copyright © 2017, Juniper Networks, Inc.

```
     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
```


Copyright © 2017, Juniper Networks, Inc.

test aaa dhcp user

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*>

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.

Description

Verify Dynamic Host Configuration Protocol (DHCP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

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.

Copyright © 2017, Juniper Networks, Inc.

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

Related Documentation

- Testing a Subscriber AAA Configuration on page 823

List of Sample Output test aaa dhcp user on page 1871

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

Copyright © 2017, Juniper Networks, Inc.

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>
```

Copyright © 2017, Juniper Networks, Inc.

```
     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
```


Copyright © 2017, Juniper Networks, Inc.

test aaa ppp user

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*>

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.

Description

Verify Point-to-Point Protocol (PPP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

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.

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

Copyright © 2017, Juniper Networks, Inc.

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

Related Documentation

- Testing a Subscriber AAA Configuration on page 823

List of Sample Output

test aaa ppp user on page 1875
test aaa ppp user (tunneled user) on page 1876
test aaa ppp user (authentication failure) on page 1876

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>**.

Copyright © 2017, Juniper Networks, Inc.

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
```

Copyright © 2017, Juniper Networks, Inc.

```
        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:

Copyright © 2017, Juniper Networks, Inc.

```
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 (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 -                          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 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
```

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.